


Rapid Optoelectronic Characterization of Semiconductors by Combining Bayesian Inference with Metropolis Sampling

Calvin Fai[✉], Charles J. Hages[✉], and Anthony J.C. Ladd^{✉*}

Department of Chemical Engineering, University of Florida, Gainesville, Florida 32611, USA

 (Received 8 June 2023; revised 6 July 2023; accepted 26 July 2023; published 1 September 2023)

Quantifying rates of charge carrier recombination is a crucial step in developing solar cells with high power conversion efficiencies. However, recovering characteristic parameters such as carrier mobility, doping concentration, and recombination rate constants is hindered by the interplay of the carrier dynamics with multiple recombination mechanisms. Interpretation of optoelectronic measurements, such as time-resolved photoluminescence (TRPL), usually relies on analytically tractable simplifications to the underlying physics models for carrier mobility and recombination, which sacrifices some of the information content of the measurement. We have recently shown that, by incorporating simulations of the complete carrier physics into a Bayesian analysis, previously unrecoverable material parameters, such as carrier mobility and the doping level, can be determined from TRPL measurements. Unfortunately, the large number of simulations required by a random sampling of the parameter space necessitated access to high-performance computing resources, limiting the usefulness of the approach. Here, we introduce an importance sampling algorithm (Metropolis Monte Carlo), which reduces the computational requirements by 2–3 orders of magnitude, rendering the Bayesian inference tractable on desktop computers. These developments affirm the utility of a simulation-driven analysis of optical characterization measurements, and make a physics-informed Bayesian inference available to all semiconductor researchers.

DOI: [10.1103/PRXEnergy.2.033013](https://doi.org/10.1103/PRXEnergy.2.033013)

I. INTRODUCTION

Research in photovoltaic cells has delivered dramatic increases in power conversion efficiency (PCE) over the past decade, with the top halide perovskite nearly doubling its PCE in that timeframe to 26% [1]. However, maturation of its precedents, such as crystalline silicon and cadmium telluride devices [2], serves as a reminder that rapid PCE growth is difficult to sustain. A major limitation to further increases in PCE is the disconnect between our ability to measure manifestations of the carrier physics governing the performance of semiconductor devices and our ability to process actionable information from these measurements. Measurements such as time-resolved photoluminescence (TRPL), terahertz spectroscopy, and transient absorption curves allow us to manipulate populations of charge carriers within a material and quantify their decay. However, underlying the observed transients is a web of

interconnected carrier transport and recombination mechanisms, which are difficult to decouple. In the case of TRPL, for example, it is common to report one or two effective carrier lifetimes calculated by a mono- or biexponential fit to a “low-injection” limiting region of the curve [3–6]. Yet, a deeper analysis shows that these lifetimes are composites of several mechanisms [7–9], such as carrier detrapping [10] and defect-assisted recombination at both bulk and surface sites [11,12]. Quantifying each of these mechanisms is necessary in order to pinpoint key areas of improvement in material design. However, the computational expense of the carrier transport model [13,14], combined with the large number of instances needed to sample a multidimensional parameter space, has precluded a comprehensive analysis of the available characterization data.

Bayesian inference approaches have demonstrated the ability to analyze complex parameter spaces and provide a more complete picture of the dominant loss mechanisms within specific materials [15,16]. In Bayesian inference, a likelihood function is created to quantify the degree of correlation between a simulated data set, determined by parameters in the physics model, and the actual experimental data. A statistical sampling is then performed on likelihood functions from many different sets of parameter values, which are combined to form a posterior

*tladd@che.ufl.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

probability distribution, in accordance with Bayes' theorem. The posterior distribution is then analyzed to obtain a prediction and confidence measure of the optimum parameter set that best describes the measurements, and reveal any covariances between different parameters. In 2017, Brandt *et al.* [15] used a Bayesian inference method with adaptive grid sampling to extract the minority carrier mobility, conduction-band offset, interface recombination velocity, and minority carrier lifetime from current-voltage measurements of SnS cells. Recently, we applied a GPU-accelerated Bayesian inference algorithm to standard TRPL measurements of single absorbers [16]. Up to ten optoelectronic parameters were recovered, describing carrier diffusion, doping concentration, and recombination through Shockley-Reed-Hall, Auger, and band-to-band radiative mechanisms. Bayesian inference can determine precise parameter values and elucidate details of the underlying physics, while at the same time proceeding with limited human intervention. Nevertheless, the automated analysis incurred a high computational cost, requiring a state-of-the-art GPU-accelerated supercomputer [16]. Here we describe an improved sampling scheme, which reduces the computational requirements by orders of magnitude, and enables a complete analysis of a typical absorber using a desktop workstation.

Metropolis Monte Carlo (MMC) sampling [17] generates a Markov sequence of states Ω_k (in our case sets of parameter values) sampled from a distribution $P(\Omega)$. In statistical physics, P is formed from the Boltzmann distribution of the Hamiltonian; here P is the product of likelihood functions from each data point. Importance sampling avoids states with low probability and is therefore much more efficient computationally than either quadrature methods or random sampling. Our goal in this paper is to present compelling evidence of the effectiveness of Bayesian inference as a practical tool for optoelectronic characterization. An analysis of a three-fluence power scan can be completed in about 5 h using a desktop workstation (eight cores). The analysis is faster, less labor intensive, and with a significantly higher information content than current analytical methods.

We first present a brief summary of the statistical analysis underpinning the use of Bayesian inference for parameter determination (Sec. II). Results presented in Sec. III show that MMC sampling of TRPL data sets recovers similar optoelectronic properties to random sampling [16], but with significantly reduced computational requirements. We discuss a series of optimizations that can further reduce the uncertainties by an order of magnitude or more, without a significant increase in computation time. The implications of these results for advances in photovoltaic device development are outlined in Sec. IV. The construction of the Metropolis sampling in the context of a Bayesian analysis of TRPL and related measurements is described in Sec. V.

II. THEORY

Simulated TRPL decay can be represented mathematically by a functional relation $\mathbf{Y}^K = \mathbf{F}^K(\Omega)$, where \mathbf{Y}^K represents TRPL intensities at a predetermined set of K time points and Ω is a set of optoelectronic parameters that characterize the material. The function F is the standard drift-diffusion model for carrier transport and recombination [13], including contributions from radiative, Auger, and Shockley-Read-Hall recombination. Defect-mediated surface recombination is included as boundary conditions on the carrier fluxes. The physics model was extensively analyzed in previous work [16], and is summarized in the Appendix.

Since experimental measurements are intrinsically stochastic, it is natural to frame the physics model in terms of probabilities. A distribution of measurements $P(\mathbf{Y}^K)$ can be related to a distribution of parameters $P(\Omega)$ by a likelihood function $L(\mathbf{Y}^K, \Omega)$,

$$P(\mathbf{Y}^K) = \int L(\mathbf{Y}^K, \Omega)P(\Omega)d\Omega. \quad (1)$$

The likelihood function is the conditional probability of a measurement \mathbf{Y}^K , given a set of parameter values Ω . While it is straightforward to predict TRPL measurements from parameter values, the inverse problem is not well posed [18]. As a result, small uncertainties in the measurements can generate artificial oscillations in the solution [18,19]. A well-posed inverse problem requires a bounded distribution of parameters $P(\Omega)$ and a regularized likelihood kernel [18], which we take as a Gaussian localization of the difference between measured and computed TRPL data,

$$L(\mathbf{Y}^K, \Omega) = \prod_{k=1}^K (2\pi\sigma_k^2)^{-3/2} \exp \left[-\frac{[Y_k - F_k(\Omega)]^2}{2\sigma_k^2} \right], \quad (2)$$

where the subscript k indicates a particular data (time) point. The extent of the regularization is controlled by the values of σ_k , which can be adjusted in accordance with the uncertainties in the experimental data and physics model (see Sec. III D 2). We recover the deterministic physics model in the limit $\sigma \rightarrow 0$.

The ‘‘posterior’’ distribution of parameter values that best accounts for the experimental data can be found from Bayes' theorem:

$$P(\Omega, \mathbf{Y}^K) = \frac{L(\mathbf{Y}^K, \Omega)P(\Omega)}{P(\mathbf{Y}^K)}. \quad (3)$$

It contains the statistically relevant information about the inferred parameters: the most likely value, the mean value, and the variance (uncertainty). Analysis of posterior distributions was described in Ref. [16].

The multidimensional integrals in Eq. (1) are most efficiently calculated by sampling rather than quadrature, replacing Eq. (1) by a summation over a discrete set of parameter values, Ω^N :

$$P(\mathbf{Y}^K) = \sum_{n=1}^N L(\mathbf{Y}^K, \Omega_n) P(\Omega_n). \quad (4)$$

The simplest strategy to determine the posterior distribution is to sample parameter values from a bounded uniform prior distribution $P(\Omega_n) = 1/N$ and compute the posterior distribution using Eqs. (3) and (4) [16]. However, since most parameters $\Omega_n \in \Omega^N$ will have a vanishingly small likelihood, random sampling is inefficient. The computational advantage of MMC (see Sec. VB) is that it samples directly from the posterior distribution (3) and therefore avoids low probability states.

III. RESULTS

We have previously shown that Bayesian inference can recover seven key optoelectronic properties from a limited number of TRPL scans [16]: the ambipolar mobility μ' , the radiative recombination rate k^* , the doping level p_0 , the combined Auger coefficient $C_n + C_p$, the nonradiative decay times, τ_n and τ_p , and the combined surface recombination rate $S_F + S_B$. The Bayesian inference exceeded previous analytic analysis of TRPL data both in the information content (number of parameters that could be recovered) and in the precision of the inferred values. Here we introduce MMC sampling to increase the computational efficiency of a Bayesian analysis of TRPL decay.

A. Simulated TRPL for a $\text{CH}_3\text{NH}_3\text{PbI}_{3-x}\text{Cl}_x$ perovskite absorber

Simulated TRPL decay curves, shown in Fig. 1, were generated at three fluences: 2.75×10^{13} , 1.92×10^{12} , and 2.12×10^{11} photons cm^{-2} . These data sets were used in lieu of experimental measurements as a means of assessing the effectiveness of the Bayesian inference. Since the exact parameter values are known in advance, we can judge precisely how well each parameter can be inferred, and then determine the relative computational efficiencies of MMC and random sampling.

Input parameters for the simulations (Table I) were taken from an analytical analysis of several different experiments [7]: optically pumped terahertz spectroscopy, TRPL, and thermoelectric power (Seebeck effect) measurements. Experimental noise in the detection system was estimated from fluctuations in the experimental data, as described in Sec. VC. The resulting TRPL decay curves, shown in Fig. 1, were used as input to the Bayesian inference. Two sample thicknesses are needed to separate the bulk and surface contributions to the recombination rates [16].

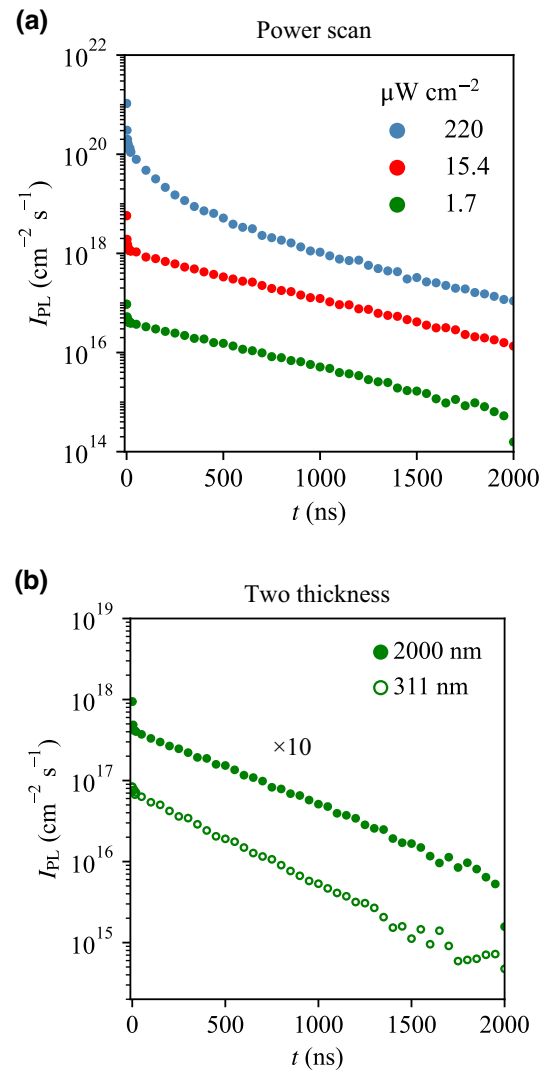


FIG. 1. Simulated TRPL curves for a perovskite absorber including a model for the experimental noise (Sec. VC). (a) Simulated three-fluence power scan for a 2000-nm sample. (b) Simulated two-thickness power scan at low injection. Two sample thicknesses are needed to separate bulk and surface recombination rates [16]. The plot for the 2000-nm sample is offset by a factor of 10 for clarity.

Marginalized probability distribution functions (PDFs) for the inferred parameters are shown in Fig. 2. The graph shown in panel (a) indicates the log likelihood [Eq. (2)] as a function of the number of steps in the Markov chain N , and indicates that optimal parameter values are located after a few thousand states (parameter sets). A discussion of the underlying physics and inferencing methodology can be found in our previous work [16]. The distribution of each parameter is approximately Gaussian, with a peak close to the input value of the parameter (red lines). Figures 2(b)–2(e) indicate the parameters that can be inferred from a single power scan with a 2000-nm sample thickness. The combined Shockley-Reed-Hall lifetime

TABLE I. Input and inferred parameters from simulated TRPL power scans; Gaussian noise [Eq. (14)] was included in the simulated intensity, with the noise levels described in Sec. VC. The inferred parameters include the ambipolar mobility μ' , the combined Auger recombination coefficient $C = C_n + C_p$, and the combined surface recombination velocity $S = S_F + S_B$. The first six rows of the table were inferred from a single three-fluence power scan on a 2000-nm sample; the remaining rows (\ddagger) included data from an additional power scan of a 311-nm sample.

Parameter	Input	Inference	
		MMC	Random [16]
μ_n^a (cm ² V ⁻¹ s ⁻¹)	20	$\mu' = 19.6 \pm 0.7$	$\mu' = 23 \pm 3.8$
μ_p^a (cm ² V ⁻¹ s ⁻¹)	20		
$k^*/10^{-11}$ (cm ³ s ⁻¹)	4.8	4.80 ± 0.04	4.75 ± 0.10
$p_0/10^{15}$ (cm ⁻³)	3	3.01 ± 0.05	3.03 ± 0.10
$C_n/10^{-29}$ (cm ⁶ s ⁻¹)	4.4	$C = 8.8 \pm 0.6$...
$C_p/10^{-29}$ (cm ⁶ s ⁻¹)	4.4		
τ_n (ns) [‡]	511	508 ± 1.4	514 ± 2.1
τ_p (ns) [‡]	871	878 ± 5.0	879 ± 9.1
S_F (cm s ⁻¹) [‡]	10	$S = 20.2 \pm 0.1$	$S = 20.7 \pm 0.1$
S_B (cm s ⁻¹) [‡]	10		

Fixed parameters: minority carrier density $n_0 = 10^8$ cm⁻³, relative permittivity $\epsilon_r = 10$, and absorption coefficient $\alpha = 6 \times 10^4$ cm⁻¹.

^aFrom short-time data, $t < 15.6$ ns.

(not shown) can also be inferred from the exponential decay of the TRPL curve in low injection. It is sometimes referred to as the minority carrier lifetime, but is in fact

a combination of bulk and surface recombination rates, $\tau_{SRH} = (\tau_n^{-1} + \tau_p^{-1} + S_F/d + S_B/d)^{-1}$. The bulk and surface recombination rates can be individually inferred by including data from an additional thickness (311 nm); the PDFs for these parameters are shown in Figs. 2(f)–2(h). In some cases only combinations of parameters can be inferred: the ambipolar mobility (μ'), the combined Auger recombination rate ($C = C_n + C_p$), and the combined surface relaxation rates ($S = S_F + S_B$). We have previously described how some of these parameters can be disentangled by additional TRPL measurements on samples with high surface recombination [16].

The parameter values inferred by MMC and random sampling are compared with the input parameters in Table I. The reported uncertainties are one standard deviation in the marginalized posterior distributions for each parameter (Fig. 2); they are somewhat smaller for MMC than for random sampling. The uncertainties are affected by the level of noise in the detector $\sigma_{E,k}$ and by the assigned model error σ (Sec. VC). We chose $\sigma = 1$ for these inferences, which is the lower limit of what can be used in the random sampling algorithm; for smaller σ , the likelihood distribution is too granular to be sampled effectively, even with 10^7 states. The effects of varying σ are addressed in Sec. III D 2, using MMC sampling only. We recalculated the inferences from random sampling [16] to include the updated uncertainty calculations.

For all data sets, the MMC method delivers a significant increase in computational efficiency (more than 200-fold) over random sampling. The MMC calculations were run on a desktop workstation, while the random sampling




FIG. 2. Parameter distributions from simulated TRPL measurements of a perovskite absorber ($\text{CH}_3\text{NH}_3\text{PbI}_{3-x}\text{Cl}_x$). (a) Trace of the log likelihood from a Markov chain generated by the MMC method. Parameters from a single-thickness power scan are the (b) ambipolar mobility μ' , (c) doping concentration p_0 , (d) radiative recombination rate k^* , and (e) Auger recombination rate $C_n + C_p$. Additionally, from a two-thickness power scan, we have (f)–(g) bulk carrier lifetimes τ_n , τ_p and (h) the surface recombination velocity $S_F + S_B$.

TABLE II. Computational time to obtain the parameter inferences reported in Table I. We report timings for MMC on a single core and random sampling on a DGX-A100 using eight GPUs simultaneously. The timings for random sampling on a single CPU are an estimate from a limited subset of states.

Test case	MMC CPU	Random GPU ($\times 8$)	Random CPU
Power scan	21.7 h	16.8 h	5680 h
	$N = 32\,000$	$N = 8.4 \times 10^6$	$N = 8.4 \times 10^6$
Two thickness	19.4 h	33.5 h	10 150 h
	$N = 16\,000$	$N = 8.4 \times 10^6$	$N = 8.4 \times 10^6$

algorithm used a DGX-A100 supercomputer, consisting of eight GPUs running in parallel. The timings are summarized in Table II; the CPU timings for random sampling are estimates based on shorter runs. As a point of comparison, a desktop computer with one Intel Core i7-4790 can simulate the 96 000 TRPL curves for the two-thickness three-fluence power scan (Table I) in 20 h using the MMC method. To achieve a similar uncertainty level, random sampling required approximately 50 million simulations, requiring 270 GPU hours (34 h on eight GPUs) or 10 000 CPU hours.

B. Inferring parameters from experimental TRPL

The long-term goal of our work is to improve the analysis of experimental TRPL and related measurements. As a first step, here we use raw TRPL data as input, to compare our inferred parameters with values from an analytic analysis of the same data set [7]. However, the experimental data sets were normalized to one at $t = 0$, whereas the simulations were normalized to the expected fluence of the laser pulse. We therefore included an amplitude scaling between simulated and experimental TRPL curves as part of the inferencing. In Fig. 3 the simulated TRPL are scaled to the laser fluence (as in Fig. 1), and the inferred amplitudes are used to map the experimental curves to the same scale.

The experimental TRPL data are limited to a single sample thickness, and only a subset of the parameters can be inferred. In particular, the sample thickness (311 nm) is insufficient to probe carrier redistribution, so the ambipolar mobility cannot be inferred from these measurements [16]. The parameters that can be inferred— k^* , p_0 , τ_{SRH} , and $C = C_n + C_p$ —are summarized in Table III. Simulated TRPL curves using our inferred parameters are compared with the original data sets in Fig. 3(a); the results confirm the overall accuracy of the inferencing.

Values of k^* and τ_{SRH} (Table III) from analytic analysis agree with Bayesian inference within two standard deviations. However, the combined Auger recombination rate from analytical analysis differs from the Bayesian




FIG. 3. Comparison of simulated and experimental TRPL data for the 311-nm absorber. (a) Simulated (lines) and experimental (circles) TRPL at three different fluences, using the inferred parameters from Bayesian Monte Carlo (Table III, MMC). (b) Comparison of high-fluence simulations using parameters inferred from Bayesian Monte Carlo (black) and analytic analysis (red) with experimental results (blue circles).

inference value by a factor of two, outside the combined statistical uncertainties. A comparison of a simulated high-injection

(220 μW) TRPL with experiment, shown in Fig. 3(b), suggests that the initial decay rate from the analytically inferred Auger rate constant [7] is too small, whereas the curve derived from Bayesian inference is in better agreement with the measured TRPL data. The discrepancy can be attributed to the inherent limitations of an analytic analysis, which must contend with overlapping contributions from diffusion, radiative recombination, and Auger recombination. A crucial advantage of Bayesian

TABLE III. Parameter inferences from experimental TRPL [7].

Parameter	Analytic analysis [7]	MMC
$k^*/10^{-11}$ (cm ³ s ⁻¹)	4.78 ± 0.43	5.90 ± 0.22
$p_0/10^{15}$ (cm ⁻³)	$\approx 3^a$	6.2 ± 2.1
τ_{SRH} HI (ns)	1380 ± 260	1585 ± 140
$C/10^{-29}$ (cm ⁶ s ⁻¹)	8.8 ± 1.6	19.0 ± 1.5

^aDetermined from thermoelectric power data.
 $C = C_n + C_p$.

inference is that the forward simulation model automatically captures coupled recombination mechanisms. Finally, Staub *et al.* [7] were not able to determine p_0 from their TRPL data, and the approximate value quoted in their paper is from thermoelectric power measurements. However, Bayesian inference was able to deduce values for p_0 from TRPL data sets with uncertainties of 35%.

C. Determining an initial state for the Markov chain

In general, states in an MMC chain will automatically drift to high probability regions of the parameter space after a sufficient number of steps. However, it is possible that the system can become locked in a local maximum of probability and then take a long time to find the global maximum. It can also happen that parameters can drift into regions where the physical model or the simulation method breaks down; this typically happens if the defect-mediated relaxation times become much too small. To prevent such states from occurring, we establish bounds on the parameters such that trial moves outside these bounds are automatically rejected. The bounds on the prior distribution are similar to the sampling volume in previous work [16] and are listed in Table IV. In practice, any starting point within (and possibly beyond) the bounded prior leads to the same equilibrium region in the sample space,

but it may still be advantageous to find a plausible initial state within these bounds; the values used in this work are shown in Table IV.

Formulas to estimate the parameter values are given in the fourth column of Table IV. From the sample thickness d and its absorption coefficient α , the initial carrier density profile $\Delta n(t=0)$ can be calculated from the laser fluence I_0 . The distribution of carriers is initially exponential, but there is a rapid relaxation, over time scales $\tau_D \sim 10$ ns, to an essentially uniform profile. This appears as an almost instantaneous drop in PL intensity (Fig. 1), and the delay time can be used to estimate the ambipolar mobility $\mu' = 2\mu_n\mu_p/(\mu_n + \mu_p)$. After the initial relaxation, nearly uniform carrier distributions decay by radiative and nonradiative recombination mechanisms [7,13,16]. Most of the remaining parameters can be estimated from three additional data points: the initial PL intensity in high injection, $I_{\text{PL}}(0) \approx 10^{21}$ cm⁻² s⁻¹, the inverse slope of the long-time decay $\tau_{\text{SRH}} = 460$ ns, and the intensity found by extrapolating the low-injection decay back to zero time, $I_{\text{PL}}^{\text{SRH}}(0) \approx 10^{17}$ cm⁻² s⁻¹ (Fig. 1). Assuming that carrier redistribution is rapid in comparison with recombination, we can derive p_0 from the extrapolated intensity of the SRH decay, $I_{\text{PL}}^{\text{SRH}}(0) = k^*I_0p_0$. Estimates of the various recombination times can be obtained from the slope of the TRPL decay at low-injection levels, τ_{SRH} . We have assumed that the initial Auger and radiative recombination rates are similar.

D. Optimizing the MMC method

The results in Table I show that the combination of Bayesian inference with MMC sampling enables accurate predictions of most optoelectronic parameters from TRPL data alone (Sec. III A). The results are of comparable precision to those obtained by random sampling, but with the computational demands reduced to a tractable level for a desktop workstation (Table II). Further improvement in computational efficiency can be obtained by optimizing

TABLE IV. Guidelines for constructing an initial state for the Markov chain. Estimates for the initial parameters are listed along with the formulas used to derive the values; here, q represents the elementary charge and I_0 is the laser fluence.

Parameter	Estimated value	Bounds	Formula
μ_n (cm ² V ⁻¹ s ⁻¹)	15.6	(1, 100)	$\mu_n = \mu_p = qd^2/\pi^2k_B T\tau_D$ [16]
μ_p (cm ² V ⁻¹ s ⁻¹)	15.6	(1, 100)	
S_F (cm s ⁻¹)	219	(10 ⁻⁴ , 10 ⁴)	$S_F = S_B = d/2\tau_{\text{SRH}}$
S_B (cm s ⁻¹)	219	(10 ⁻⁴ , 10 ⁴)	
k^* (cm ³ s ⁻¹)	4.4×10^{-11}	(10 ⁻¹¹ , 10 ⁻⁹)	$I_{\text{PL}}(0) = \frac{1}{2}k^*I_0^2\alpha$
p_0 (cm ⁻³)	8.6×10^{15}	(10 ¹⁴ , 10 ¹⁶)	$I_{\text{PL}}^{\text{SRH}}(0) = k^*I_0p_0$
C_n (cm ⁶ s ⁻¹)	2.0×10^{-29}	(10 ⁻²⁹ , 10 ⁻²⁷)	$I_{\text{PL}}^a(0) = \frac{1}{3}(C_n + C_p)I_0^3\alpha^2$
C_p (cm ⁶ s ⁻¹)	2.0×10^{-29}	(10 ⁻²⁹ , 10 ⁻²⁷)	
τ_n (ns)	456	(1, 1500)	$\tau_n = \tau_p = \tau_{\text{SRH}}$
τ_p (ns)	456	(1, 3000)	

^aAssuming that the Auger and radiative recombination rates are roughly equal.




FIG. 4. Effect of trial move displacement on the rate of equilibration. (a) Traces showing the evolution of Markov chains towards the equilibrium (most probable) states for $\sigma = 1$: blue, $\Delta = 10^{-2}$; orange, $\Delta = 10^{-3}$; green, $\Delta = 10^{-4}$; red, $\Delta = 10^{-5}$. (b) Mean acceptance rate f (black squares) and state space diffusion constant D_Ω (circles) as a function of the scaled step size Δ/σ . Acceptance rates are independent of σ ($\sigma = 1$ is shown), while the diffusion scales as σ^2 .

the choice of hyperparameters: Δ , the proposal distribution width, and σ , the model uncertainty. In the language of machine learning, a hyperparameter is one that is used to control the learning process, as opposed to one of the parameters that is acquired from the data. A summary of the MMC method and the role of these hyperparameters is provided in Sec. V.

In contrast to random sampling, which distributes likelihood calculations over all regions of the prior distribution, the evolution of a Markov chain is sensitive to the hyperparameters σ_k^2 appearing in the likelihood function [Eq. (8)]. Small σ_k results in narrow distributions that the Markov process may find difficult to locate, while a large σ_k results in a broader distribution with greater uncertainties in the inferred parameters. We performed a sensitivity analysis on the rate of equilibration as a function of a constant uncertainty hyperparameter $\sigma_k = \sigma$, combining the experimental and model uncertainties into a single value for σ . The equilibration period, which is the number of states the MMC samples between its initial state and the stationary distribution was used to determine the optimum Δ . To quantify the accuracy of the inference, we calculate the mean and variance of the postequilibration samples and compare them against the known input values from Table I.

1. Selecting the trial move displacement (Δ)

A single step in MMC consists of a trial move to a proposed state, followed by acceptance or rejection of the move based on the relative likelihoods of the current and proposed states. In our implementation, states Ω_n represent the logarithm of the parameter values, so that a small change in Ω_n describes a relative change in the parameter.

Using logarithms brings the variations in disparate parameters to similar scales. Trial moves are selected from a joint uniform distribution of width Δ centered on the current state Ω_n ,

$$\Omega'_n = \Omega_n + \Delta \mathcal{R}, \quad (5)$$

where \mathcal{R} is a vector of random numbers, each one uniformly distributed in the range $[-1, 1]$. A trial move that lies outside the bounds of the prior distribution $\Omega'_n \notin \Omega^P$ is always rejected. Repeated detection of trial moves attempting to leave the bounded region is an indication that the bounds need to be increased. We do not observe any attempted crossings for the initial parameters and bounds shown in Table IV.

The magnitude of Δ exerts a significant influence on the equilibration time [Fig. 4(a)], so it can be regarded as an optimizable hyperparameter. Intuitively, a small Δ will result in slow convergence toward the stationary distribution, while a large Δ may not have sufficient granularity to locate small regions of high likelihood. Convergence difficulties can be recognized by monitoring the acceptance rate, f , the proportion of proposed moves that lead to accepted states, averaged over the whole Markov chain. A quantitative measure indicating a well chosen Δ is a maximum in the “diffusion constant” in the state space [20],

$$D_\Omega = \langle H \mathcal{R}^2 \rangle \Delta^2 < \frac{f \Delta^2}{3}; \quad (6)$$

here $H = 1$ for accepted moves and zero otherwise. The angle brackets indicate an average over the Markov chain and the upper bound follows from decomposing the correlated average $\langle H \mathcal{R}^2 \rangle$ into the individual averages, $\langle H \rangle =$

f and $\langle \mathcal{R}^2 \rangle = 1/3$; the upper bound to D_Ω [Eq. (6)] is within 20% of the actual value.

MMC simulations of the three-fluence power scan (Sec. III A), with constant uncertainty $\sigma_k = \sigma$, were used to investigate the rate of equilibration of the Markov chains and its dependence on Δ . Figure 4(b) shows that the optimum trial move displacement (Δ) is proportional to the width of the likelihood function σ . Steps larger than 0.01σ are never accepted, so the chain will remain trapped at its initial location. At the other extreme ($\Delta < 10^{-4}\sigma$), all trial moves are accepted, but equilibration becomes very slow. Rapid equilibration (< 2000 steps) is obtained with a range of Δ values $10^{-3}\sigma < \Delta < 10^{-2}\sigma$, corresponding to acceptance ratios between 0.3% and 55%. The optimum exploration of the state space (maximum D_Ω) occurs for $\Delta = 0.003\sigma$, when $f \approx 15\%$ is close to the theoretical optimum for a multivariate Gaussian posterior (23%) [21]. The relatively small changes in parameter values transitioning from the initial to equilibrium states, compared with the huge change in likelihood, roughly 10^{10000} [Fig. 2(a)], suggests that the method is sensitive to small variations in the parameters.

Further optimization may be possible with more complex algorithms for the trial moves. For instance, the differing uncertainties for each material parameter shows that the stationary distribution varies in width along each parameter. A uniform hyperrectangular distribution, where a separate Δ governs the moves for each parameter, would provide more options for tuning. Adaptive MMC algorithms [22–25] allow more efficient sampling by dynamically adjusting the shape of the trial move distribution as the Markov chain develops.

2. Effect of model uncertainty on the parameter inference

The model uncertainty σ has the effect of broadening a distribution that may otherwise be too granular. On the other hand, a large σ can significantly increase the uncertainty in the inference and lead to a displacement of a marginalized distribution from the true mean (see Fig. 2). The qualitative effect of σ on the sampling can be understood from traces through a two-dimensional parameter space of k^* and p_0 , which are shown in Fig. 5. While all random walks sample the desired target ($p_0 = 3 \times 10^{15} \text{ cm}^{-3}$, $k^* = 4.8 \times 10^{-11} \text{ cm}^3 \text{ s}^{-1}$), for the largest σ , there is significant sampling of the skewed Gaussian likelihood surrounding the target values. This results in an unnecessarily broad distribution and can result in distorted marginalizations when sampling higher-dimensional spaces.

To optimize the choice of σ , we simulated TRPL curves without experimental noise and ran the inference with varying values of σ . Results reported in Table V show that the inference is robust to a wide variation in




FIG. 5. MMC sampling of k^* and p_0 using different values of σ . The input (correct) values of k^* and p_0 are indicated with black dashed lines.

model σ . Acceptable accuracy can be obtained from broad likelihood distributions ($\sigma = 10$), while essentially perfect inferences can be found for $\sigma = 1$, which was the value used in Sec. III A. Even smaller values of σ do not introduce any apparent lack of ergodicity or disconnected regions of the sample space. The inferred parameter values are within one standard deviation of the input values for all $\sigma > 0.01$. At $\sigma = 0.01$ the errors in the inferred parameters are very small, less than 0.01%, but occasionally more than two standard deviations away from the input value.

Results in Table VI include the same experimental noise [Eq. (14)] as in Table I. The experimental noise was included in the simulated TRPL data (Fig. 1) and in the likelihood width σ_k [Eq. (9)]; the maximum trial displacement was set to $3 \times 10^{-3}\sigma$, as in the noise-free case. The results show that realistic levels of experimental noise do not significantly harm the inferencing; the inferred values are within 1% of the input value, except when the noise is dominated by model uncertainty ($\sigma = 10$). Parameters that control the decay at high-injection levels (k^* , p_0 , C , S) do not see the effects of the background noise and their uncertainty scales as $\sqrt{\sigma^2 + \sigma_1^2}$. However, parameters that control decay at low-injection levels (τ_n , τ_p) have additional uncertainty due to the background noise σ_0/I_k . At low-injection levels TRPL decay is controlled by τ_n , the minority carrier lifetime. Background noise will tend to raise the detected signal, leading to a larger estimate of τ_n and a smaller value for τ_p , because the bulk Shockley-Reed-Hall decay at high-injection levels is linked to the sum $\tau_n + \tau_p$, making for a small cross-correlation between

TABLE V. Effect of σ on parameter predictions from a noise-free two-thickness power scan. Quantities are defined in the same units as Table I.

Parameter	Input	$\sigma = 10$	$\sigma = 1$	$\sigma = 0.1$	$\sigma = 0.01$
$k^*/10^{-11}$	4.8	4.79 ± 0.18	4.797 ± 0.021	4.799 ± 0.002	4.7991 ± 0.0002
$p_0/10^{15}$	3	3.03 ± 0.25	3.005 ± 0.035	3.003 ± 0.003	3.0006 ± 0.0003
$C/10^{-29}$	8.8	8.92 ± 1.02	8.77 ± 0.14	8.790 ± 0.010	8.802 ± 0.001
τ_n	511	512 ± 14	511.3 ± 1.6	511.01 ± 0.14	511.02 ± 0.02
τ_p	871	865 ± 47	870.4 ± 5.1	870.88 ± 0.56	871.03 ± 0.05
S	20	19.8 ± 1.3	20.01 ± 0.10	20.001 ± 0.013	20.005 ± 0.001
f	...	0.143	0.162	0.153	0.154

$$C = C_n + C_p; S = S_F + S_B.$$

the two relaxation times. However, the inferred τ_n , τ_p remain within 1% of the input values. We also note that the inferences of S can be outside the estimated error bars as well. For this material, surface recombination is always more than an order of magnitude slower than bulk recombination, making the inference more prone to experimental errors. We note that in the noise-free case all these parameters can be inferred essentially exactly.

E. Parallel sampling from different initial states

A drawback of the MMC algorithm is that it is inherently serial. Unlike random sampling, where many states (parameter sets) can be considered simultaneously [16], a Markov chain can only generate one new state at a time. A degree of parallelism can be achieved by combining shorter chains from a number of different initial conditions. This has two advantages: by varying the initial state, it may be possible to identify situations where a Markov chain is trapped in a limited region of the parameter space, and the statistics from a number of parallel simulations can be combined. However, since each simulation requires N_0 equilibration steps before it can start to accumulate statistics, this sets an upper limit on the number of parallel threads that can be effectively utilized of about N/N_0 . An additional parallelism can be obtained by distributing the simulation domain (the thickness of the sample) across

multiple threads. However, this is also somewhat limited by the computational cost of the boundary nodes, which cannot be effectively parallelized. Overall, it might be possible to usefully run up to 100 threads in parallel, much less than the tens of thousands available on GPU-based computers, but still a significant enhancement over serial processing.

A limited parallelism was explored by repeating the inference of Sec. III A with eight independent walks, initiated within the bounds on the parameters indicated in Table IV. Traces of selected parameters (p_0 , k^* , τ_p) for the different Markov chains are illustrated in Fig. 6 along with the log likelihood function. The different initial states rapidly locate to the same high-likelihood region of the parameter space and then oscillate around the expected values after the same number of equilibration steps, $N_0 \approx 2000$. The results confirm that the most probable region of the parameter space can be accessed from a wide range of initial conditions, with the initial values of individual parameters varying by two orders of magnitude. We illustrate the performance improvement in Table VII by determining the run time for different numbers of chains, each with a cumulative total of 30 000 states. The workload is divided evenly among the M CPUs, with an overhead due to the equilibration common to all the Markov chains. An eight-core simulation is five times faster than a single core, for a parallelization efficiency of about 60%.

TABLE VI. Effect of σ on parameter predictions from a two-thickness power scan with Gaussian noise [Eq. (14)]; the noise levels are given in Sec. V C. Quantities are defined in the same units as Table I.

Parameter	Input	$\sigma = 10$	$\sigma = 1$	$\sigma = 0.1$	$\sigma = 0.01$
$k^*/10^{-11}$	4.8	4.81 ± 0.24	4.820 ± 0.023	4.786 ± 0.002	4.7940 ± 0.0006
$p_0/10^{15}$	3	3.00 ± 0.35	3.00 ± 0.034	2.986 ± 0.004	2.9890 ± 0.0009
$C/10^{-29}$	8.8	9.9 ± 1.6	8.75 ± 0.14	8.842 ± 0.014	8.840 ± 0.004
τ_n	511	508 ± 14	508.5 ± 1.4	514.10 ± 0.18	514.31 ± 0.08
τ_p	871	898 ± 48	875.0 ± 4.7	861.94 ± 0.55	863.98 ± 0.19
S	20	20.7 ± 1.3	20.20 ± 0.12	19.704 ± 0.017	19.829 ± 0.005
f	...	0.161	0.155	0.211	0.615

$$C = C_n + C_p; S = S_F + S_B.$$




FIG. 6. Traces from an ensemble of eight different Markov chains, with starting states at the bounds (Table IV) of three selected parameters: p_0 , k^* , and τ_p . (a) Log likelihood for states along the Markov chains. (b)–(d) Traces for the doping concentration p_0 , radiative recombination rate k^* , and majority carrier lifetime τ_p .

IV. DISCUSSION

In this work, we have introduced MMC sampling to the recovery of optoelectronic parameters from TRPL and related measurements. A suite of test cases shows that the

TABLE VII. Inferences from a single-thickness power scan with varying numbers of simultaneously generated Markov chains. The total number of equilibrated states sampled is the same in all cases ($N = 30\,000$).

# chains	N/walk	Runtime (h)	Equilibration (N_0/N) ^a (%)
1	32 000	21	6
2	17 000	12	11
4	9500	7.5	21
8	5750	4.6	35

^a $N_0 = 2000$ for all chains.

increased sampling efficiency of MMC yields competitive predictions with a GPU-accelerated random sampling but with significantly reduced hardware requirements. Optimizations of the MMC method (Sec. III D), combined with modest computational resources (an eight-core workstation), allow for a complete Bayesian inference (two sample thicknesses with three different injection levels for each sample) in about 6 h, which is five times faster than the eight-GPU DGX-A100 supercomputer using random sampling.

In Sec. III we simulated TRPL datasets using parameters derived from a variety of optoelectronic measurements [7]. Our results showed that in the absence of experimental noise we can recover the input parameters essentially exactly, to within 0.01%. For parameters that control high-injection decay (k^* , p_0 , C), the addition of experimental noise has a negligible effect on the inference, but parameters influencing low-injection decay (τ_n) can develop

cross-correlations due to the background noise. Nevertheless, these parameters can still be inferred with uncertainties of less than 1%, which is more precise than a typical analytic analysis.

Both the analytic analysis and the Bayesian inference assumed that the laser fluence was entirely absorbed by carrier excitation, but the absolute number of excited carriers is not measurable in a transient setup. Variations in fluence affect both the overall amplitude of the PL and the shape of the transient decay [Fig. 3(a)]. Thus, the actual uncertainty in the parameters may be larger than indicated by variances in the inferencing. Work is in progress to characterize the sensitivity of the inference to uncertainties in the initial carrier density.

Our results highlight the potential of Bayesian Monte Carlo to explore coupled physical phenomena that are difficult to separate analytically. Our results show that Bayesian inference can pinpoint the most significant recombination mechanisms from a small number of TRPL scans. This information can then be used to guide further treatment of the material to increase the photovoltaic conversion efficiency. Bayesian inference of TRPL data can also be used to quantify carrier trapping and detrapping [10,26], analyze carrier dynamics in different device architectures [10,26,27], and explore carrier and photon transport in nanostructures [28]. Moreover, the methodology is not limited to TRPL. For instance, the photoconductivity decay measured in an optical-pump terahertz probe (OPTP) also originates from photoexcited carriers [29–32]. With a few modifications to the physical model, Bayesian Monte Carlo can be adapted to extract fast carrier diffusion and surface recombination from OPTP data, analyze interlayer transport and exciton dynamics from transient absorption [33], or obtain complementary information from combinations of different experimental techniques. For materials characterized by more rapid carrier recombination, such as CdTe, the instrument response can be convolved with the simulated TRPL data. Our work demonstrates that the introduction of MMC sampling to Bayesian inference can significantly improve the information retrieval from material characterization experiments for a broad range of photovoltaic materials.

Finally, we note that there are a number of indicators suggesting that the parameter space is very smooth in comparison with typical Monte Carlo simulations in molecular or statistical physics. In Fig. 5 the region of exploration shrinks smoothly as σ^2 (equivalent to temperature) is reduced by orders of magnitude, without any indication that the Markov chains become trapped in separate regions of the parameter space, as happens in systems exhibiting phase transitions. Moreover, Fig. 6 shows that we quickly (within 2000 trial moves) reach the same equilibrium state from widely (two orders of magnitude) different starting states.

V. METHODS

The results shown in Sec. III were obtained from TRPL data sets, using Bayesian inference to obtain distributions of likely values of the material parameters, denoted collectively as Ω (see Table I). For a specific choice of parameters Ω_n , we simulate the expected experimental outcomes from K measurements (for example, a TRPL power scan with K total data points) $\mathbf{F}^K(\Omega_n)$ using standard device simulation models [13,16]. We assume that the experimental outcomes, denoted collectively by \mathbf{Y}^K , deviate from the predictions of a given parameter set $\mathbf{F}^K(\Omega_n)$ by an error term δ^K ,

$$Y_k = F_k(\Omega_n) + \delta_k. \quad (7)$$

Here Y_k and F_k indicate the logarithm of the TRPL intensity; we compare logarithms to avoid overly biasing the inferencing toward early delay times, as TRPL decay typically spans several orders of magnitude. The uncertainty in an experimental measurement then becomes a relative uncertainty in the logarithm. Similarly, we use logarithms of the material properties as parameters (Ω) so that each property is sampled over a similar range, relative to its value. We assume that δ^K is Gaussian distributed, with zero mean and standard deviations σ^K , but systematic errors in the experimental measurements and inadequacies in the underlying physics model can shift and skew the distribution. A strength of Bayesian inferencing methods is that they offer the possibility of a post-hoc assessment of the effectiveness of the model, by examining the distributions of inferred parameters.

A. Likelihood function

For a Gaussian process model, the log likelihood of a particular set of parameter values Ω_n is the sum of deviations between the measured data set \mathbf{Y}^K and the corresponding simulated data $\mathbf{F}^K(\Omega_n)$,

$$\ln L(\Omega_n, \mathbf{Y}^K) = - \sum_{k=1}^K \left(\frac{(Y_k - F_k(\Omega_n))^2}{2\sigma_k^2} \right). \quad (8)$$

The normalizations in Eq. (2) can be omitted since we only require ratios of likelihood functions. The parameter σ_k , representing the variance in δ_k , contains two components: the model variance σ and the uncertainty in the experimental measurements $\sigma_{E,k}$,

$$\sigma_k^2 = (\sigma^2 + \sigma_{E,k}^2). \quad (9)$$

The model uncertainty σ is a hyperparameter as it is external to the inference itself, while the $\sigma_{E,k}^2$ are determined from the experimental data sets (Sec. VC). The development of criteria for selecting a suitable value for σ is discussed in Sec. III D 2.

For a collection of N parameter sets Ω^N , sampled uniformly within a bounded space Ω^P , the posterior distribution is simply

$$P(\Omega_n, \mathbf{Y}^K) = \frac{L(\Omega_n, \mathbf{Y}^K)}{\sum_{n'=1}^N L(\Omega_{n'}, \mathbf{Y}^K)}. \quad (10)$$

A bounded prior distribution is useful in excluding parameters where the physics model or its numerical solution breaks down. The parameter ranges used in the inferencing are indicated by the ranges shown in Table IV. Parameters outside the preimposed bounds are not considered in the inferencing.

Sampling from a uniform distribution [Eq. (10)] becomes inefficient as the dimensionality of the parameter space increases because the overwhelming majority of the samples will have essentially zero probability. However, by using importance sampling, where a state Ω_n is sampled from the desired posterior probability $P(\Omega_n, \mathbf{Y}^K)$, regions of low probability are avoided and the sampling efficiency can be greatly increased. This can be accomplished by the Metropolis Monte Carlo method [17].

B. Metropolis Monte Carlo

The problem with sampling from the posterior distribution [Eq. (10)] is that the normalization is unknown *a priori*. The Metropolis algorithm [17] neatly avoids this difficulty by arranging for transitions between successive states in the Markov chain Ω_n and Ω_{n+1} to occur with relative probabilities $\pi_{n \rightarrow n+1}$ and $\pi_{n+1 \rightarrow n}$ given by

$$\frac{\pi_{n \rightarrow n+1}}{\pi_{n+1 \rightarrow n}} = \frac{P_{n+1}}{P_n} = \frac{L_{n+1}}{L_n}, \quad (11)$$

where the normalization of P in Eq. (10) cancels out. Starting from an initial state Ω_0 , with $P(\Omega_0) \neq 0$ (Sec. III C), a Markov chain of states can be generated by rejection sampling, with transition probabilities satisfying Eq. (11). For a sufficiently long chain, this enforces detailed balance between any pair of states $P_n \pi_{n \rightarrow n'} = P_{n'} \pi_{n' \rightarrow n}$, which is a sufficient condition for a stationary distribution [17].

Metropolis sampling is based on the observation that it is not possible to sample directly from the likelihood function. We must first generate a small trial displacement within the parameter space $\Omega'_n = \Omega_n + \Delta_n$, where the displacement Δ_n is sampled from a bounded uniform distribution (see Sec. III D 1). The trial move is accepted ($\Omega_{n+1} = \Omega'_n$) with probability $P_{n \rightarrow n+1}^A$ and rejected ($\Omega_{n+1} = \Omega_n$) with probability $1 - P_{n \rightarrow n+1}^A$. Since the trial move probabilities for $\Omega_n \rightarrow \Omega_{n+1}$ and $\Omega_{n+1} \rightarrow \Omega_n$ are equal, the ratio of acceptance probabilities $P_{n \rightarrow n+1}^A / P_{n+1 \rightarrow n}^A$ must

follow the ratio of the likelihood functions:

$$P_{n \rightarrow n+1}^A = \min\left(\frac{L_{n+1}}{L_n}, 1\right). \quad (12)$$

However, it is not essential that the trial move probability is symmetric between the two states n and $n+1$; if it is not then the acceptance probability must be suitably modified [34,35].

Random walks are initiated from a state in Ω^P , under the assumption that all states (in Ω^P) are accessible from any initial condition (ergodicity). We cannot discount the possibility that there are separate regions of the parameter space that are mutually inaccessible, but we have not observed any such behavior. Typically, a random walk drifts systematically towards a stationary (or equilibrium) distribution sampled from the posterior $P(\Omega, \mathbf{Y}^K)$, with increasing fidelity as the number of steps (N) increases. The selection of N follows from a trade-off between the statistical precision of the sampling and the computational

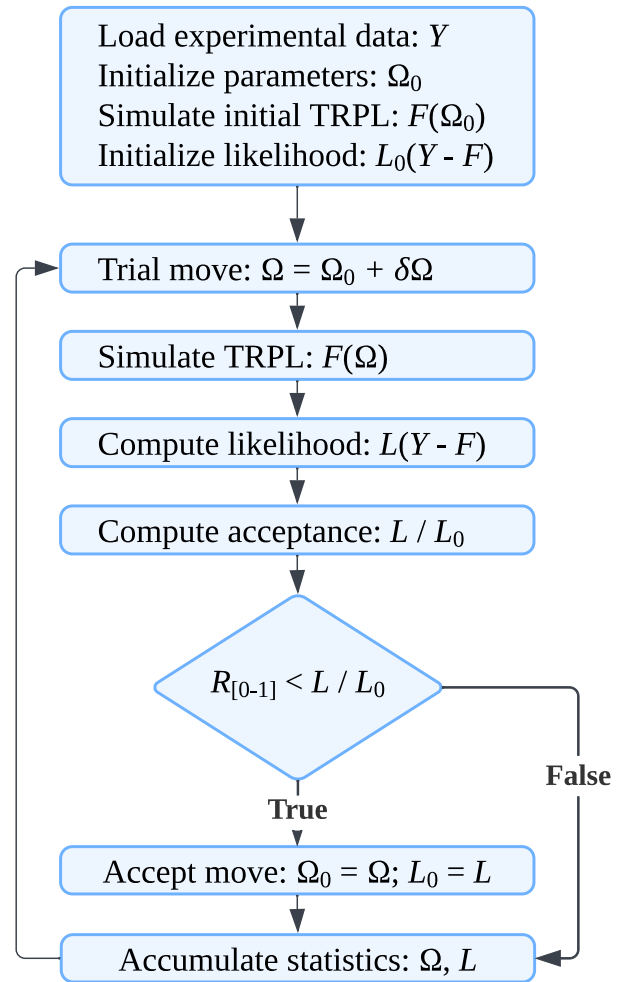


FIG. 7. A flowchart describing the Bayesian Monte Carlo algorithm.

cost; typically, N is in the range 10^4 – 10^5 . A flowchart of the MMC algorithm is shown in Fig. 7.

Several qualitative methods can be employed to search for signs of a stationary distribution. A trace of the log-likelihood from different initial states [Fig. 6(a)] reveals whether or not a stationary state has been reached and, if so, how many steps it took. Convergence to the same likelihood indicates that the random walk is oscillating between states with similarly high probability, which collectively make up a stationary distribution. There is a clear cutoff, $N_0 \approx 2000$, separating the equilibration period, in which the random walk is locating the stationary distribution, from the remaining states where the walk is oscillating within a narrow stationary distribution. Trace plots of individual parameter values display a similar convergence as the log likelihood [Figs. 6(b)–6(d)]. Averaging over the stationary portion of the random walk ($N > N_0$), we find that the parameters are Gaussian distributed and centered on the input values [Figs. 2(b)–2(h)].

The stationary distribution is not necessarily narrow for *all* parameters involved in TRPL measurements. Failure to infer a particular parameter may mean that the TRPL response is insensitive to that parameter or, more likely, that there is a strong parameter covariance, which can be revealed by higher-dimensional marginalizations [16]. Instances of parameter covariance can highlight specific aspects of the carrier physics that are not being probed by existing measurements and guide the design of more informative experiments.

C. Including uncertainty

Inferences for each parameter are made by taking mean values of the states sampled after equilibration ($N > N_0$), while the uncertainties represent one standard deviation in the distribution. Although states in the Markov chain are not independent of one another, the large number of states means that uncertainties in sampling the distribution are much smaller than the experimental and model uncertainties. We treat the model uncertainty σ as an adjustable parameter, balancing precision in the inferred value with robustness in the overall inference, meaning that the Markov chain does not get locked into local maxima in the likelihood distribution $L(\Omega)$ [Eq. (8)].

The model uncertainty σ [Eq. (9)] is a hyperparameter, which was determined empirically (Sec. III D 2). Based on measured TRPL data we decompose the experimental noise into a constant background and a contribution proportional to the intensity:

$$\sigma_{I,k}^2 = \sigma_0^2 + \sigma_1^2 I_k^2. \quad (13)$$

The proportional contribution to the noise $\sigma_1 I_k$ is dominant at high-injection levels and comes from the amplification

circuitry in the photomultiplier tube; the constant background noise σ_0 can be measured from the dark spectrum. The standard deviation in the experimental measurements is then represented by $\sigma_{I,k}$, which we parameterize by σ_0 , σ_1 , and the measured intensity I_k .

Each of the three TRPL curves from Staub *et al.* [7] was fitted by a sum of three exponentials, and the residual at each data point was taken as the difference from the fitted mean value. The parameters σ_0 and σ_1 were determined for each fluence by fitting the observed fluctuations in I_{PL} to a Gaussian distribution, with a variance given by Eq. (9). The background (low-injection) noise σ_0 was approximately $5 \times 10^{15} \text{ cm}^{-2} \text{ s}^{-1}$, $10^{15} \text{ cm}^{-2} \text{ s}^{-1}$, and $1.5 \times 10^{14} \text{ cm}^{-2} \text{ s}^{-1}$ for excitations with fluences of 2.75×10^{13} , 1.92×10^{12} , and $2.12 \times 10^{11} \text{ photons cm}^{-2}$, respectively. The amplification noise was roughly 5% of the signal ($\sigma_1 \approx 0.05$) at the highest fluence and 3% ($\sigma_1 \approx 0.03$) at the two lower fluences.

Since the likelihood function [Eq. (8)] is based on the logarithm of the measured (Y_k) and simulated (F_k) intensities, the experimental uncertainty is relative, $\sigma_{E,k} = \sigma_{I,k}/I_k$. Combining the model and experimental uncertainties from Eqs. (9) and (13), the uncertainty entering the log likelihood function [Eq. (8)] is

$$\sigma_k^2 = \sigma^2 + \frac{\sigma_0^2}{I_k^2} + \sigma_1^2. \quad (14)$$

Gaussian noise, with a standard deviation σ_k was added to the simulated TRPL decay curves (Fig. 1), and included in the σ_k parameterizing the width of the likelihood function.

The codes and input files needed to perform all analyses detailed within this study are permanently archived online [36]. The latest version of the codes are available as a public Github repository [37].

ACKNOWLEDGMENTS

The authors would like to thank Dr. Thomas Unold (Helmholtz-Zentrum Berlin) for providing published experimental data for the perovskite absorber described in Sec. IV. C.J.H. acknowledges financial support from Sony Corporation under the Sony Faculty Innovation Award program.

A.J.C.L. conceptualized the methodology; C.F. designed the codes and acquired the data. All authors contributed to analyzing the data and preparing the manuscript. A.J.C.L. and C.J.H. supervised the project; C.J.H. acquired the funding for it.

The authors declare no competing interests.

APPENDIX: PHYSICS MODEL

Charge carrier transport is based on the standard model [13] for one-dimensional carrier drift diffusion and Poisson's equation for the electric field:

$$\partial_t n = \frac{1}{q} \partial_z J_n - R, \quad (\text{A1})$$

$$\partial_t p = -\frac{1}{q} \partial_z J_p - R, \quad (\text{A2})$$

$$\partial_z E = \frac{q}{\epsilon} (\Delta p - \Delta n), \quad (\text{A3})$$

with

$$J_n = \mu_n q E n + \mu_n k T \partial_z n, \quad (\text{A4})$$

$$J_p = \mu_p q E p - \mu_p k T \partial_z p. \quad (\text{A5})$$

Here n and p are the densities of mobile electrons and holes, J_n and J_p are their charge currents, E is the electric field strength, q is the elementary charge, ϵ is the dielectric constant, and μ_n, μ_p are the electron and hole mobilities. The bulk recombination rate $R = R_{\text{Rad}} + R_{\text{Aug}} + R_{\text{SRH}}$ is the sum of contributions from radiative, Auger, and nonradiative (Shockley-Read-Hall) recombination:

$$R_{\text{Rad}} = k^* (np - n_0 p_0), \quad (\text{A6})$$

$$R_{\text{Aug}} = (C_n n + C_p p) (np - n_0 p_0), \quad (\text{A7})$$

$$R_{\text{SRH}} = \frac{np - n_0 p_0}{\tau_p n + \tau_n p}. \quad (\text{A8})$$

Carrier recombination is characterized by five inferable parameters: the *external* radiative recombination coefficient k^* (which includes the effects of photon recycling), the Auger recombination coefficients C_n, C_p , and the electron and hole nonradiative lifetimes τ_n, τ_p ; n_0, p_0 are the equilibrium free-carrier densities.

The transport equations are closed by boundary conditions on the currents:

$$J_n(z=0) = -J_p(z=0) = qR_S, \quad (\text{A9})$$

$$J_n(z=d) = -J_p(z=d) = -qR_S. \quad (\text{A10})$$

The nonradiative SRH surface recombination rate

$$R_S = \frac{np - n_0 p_0}{n/S + p/S} \quad (\text{A11})$$

can have different recombination velocities $S = S_F$ or S_B at the front and back surfaces.

The photoluminescence intensity measured in TRPL, $I_{\text{PL}}(t)$, is calculated from the radiative recombination rate,

$$I_{\text{PL}}(t) = \int_0^d R_{\text{Rad}} dz. \quad (\text{A12})$$

Initial conditions for the carrier densities are obtained from a Beer-Lambert absorption profile characterized by the laser photon fluence I_0 and the wavelength λ of the laser, which determines the absorption coefficient of the material $\alpha(\lambda)$,

$$\Delta n(t=0) = \Delta p(t=0) = I_0 \alpha e^{-\alpha z}. \quad (\text{A13})$$

A detailed description of the implementation and analysis of this physics model can be found in our previous work [16].

-
- [1] H. Min, D. Y. Lee, J. Kim, G. Kim, K. S. Lee, J. Kim, M. J. Paik, Y. K. Kim, K. S. Kim, M. G. Kim, T. J. Shin, and S. Il Seok, Perovskite solar cells with atomically coherent interlayers on SnO₂ electrodes, *Nature* **598**, 444 (2021).
 - [2] N. Kopidakis, Best research-cell efficiencies (2022), <https://www.nrel.gov/pv/assets/pdfs/best-research-cell-efficiencies-rev220630.pdf>.
 - [3] X. Huang, Y. Matsushita, H.-T. Sun, and N. Shirahata, Impact of bismuth-doping on enhanced radiative recombination in lead-free double-perovskite nanocrystals, *Nanoscale Adv.* **4**, 3091 (2022).
 - [4] R. Hertwig, S. Nishiwaki, A. N. Tiwari, and R. Carron, Investigation and mitigation of sputter damage on co-evaporated Cu(In,Ga)Se₂ absorbers for photovoltaic applications, *Sol. RRL* **6**, 2200268 (2022).
 - [5] S. Yan, J. Liu, Y. Zhou, X. Sun, Y. Zhong, X. Chen, Y. Tang, X. Guo, Q. Sun, and H. Yang, Improved minority carrier lifetime in *p*-type GaN by suppressing the non-radiative recombination process, *Appl. Phys. Express* **15**, 075501 (2022).
 - [6] L. Ye, X. Xu, S. He, Y. Liu, Y. Jin, Y. M. Yang, and H. Zhu, Molecular triplet sensitization of monolayer semiconductors in 2D organic/inorganic hybrid heterostructures, *ACS Nano* **16**, 12532 (2022).
 - [7] F. Staub, H. Hempel, J.-C. Hebig, J. Mock, U. W. Paetzold, U. Rau, T. Unold, and T. Kirchartz, Beyond Bulk Lifetimes: Insights into Lead Halide Perovskite Films from Time-Resolved Photoluminescence, *Phys. Rev. Appl.* **6**, 044017 (2016).
 - [8] L. Krückemeier, B. Krogmeier, Z. Liu, U. Rau, and T. Kirchartz, Understanding transient photoluminescence in halide perovskite layer stacks and solar cells, *Adv. Energy Mater.* **11**, 2003489 (2021).
 - [9] D. Kuciauskas, J. Moseley, and C. Lee, Identification of recombination losses in CdSe/CdTe solar cells from spectroscopic and microscopic time-resolved photoluminescence, *Sol. RRL* **5**, 2000775 (2021).
 - [10] C. J. Hages, A. Redinger, S. Levchenko, H. Hempel, M. J. Koeper, R. Agrawal, D. Greiner, C. A. Kaufmann, and T.

- Unold, Identifying the real minority carrier lifetime in non-ideal semiconductors: A case study of kesterite materials, *Adv. Energy Mater.* **7**, 1700167 (2017).
- [11] D. Kuciauskas, A. Kanevce, J. M. Burst, J. N. Duenow, R. Dhere, D. S. Albin, D. H. Levi, and R. K. Ahrenkiel, Minority carrier lifetime analysis in the bulk of thin-film absorbers using subbandgap (two-photon) excitation, *IEEE J. Photovolt.* **3**, 1319 (2013).
- [12] T. P. Weiss, B. Bissig, T. Feurer, R. Carron, S. Buecheler, and A. N. Tiwari, Bulk and surface recombination properties in thin film semiconductors with different surface treatments from time-resolved photoluminescence measurements, *Sci. Rep.* **9**, 5385 (2019).
- [13] R. Ahrenkiel, in *Minority Carriers in III-V Semiconductors: Physics and Applications*, Semiconductors and Semimetals, Vol. 39, edited by R. Ahrenkiel and M. Lundstrom (Academic Press, Inc., San Diego, 1993), p. 40, Chapter 2.
- [14] S. M. Sze, *Physics of Semiconductor Devices* (Wiley, New York, 1981), 2nd ed., p. 50, Chapter 1.7.
- [15] R. E. Brandt, R. C. Kurchin, V. Steinmann, D. Kitchaev, C. Roat, S. Levenceno, G. Ceder, T. Unold, and T. Buonassisi, Rapid photovoltaic device characterization through Bayesian parameter estimation, *Joule* **1**, 843 (2017).
- [16] C. Fai, A. J. Ladd, and C. J. Hages, Machine learning for enhanced semiconductor characterization from time-resolved photoluminescence, *Joule* **6**, 2585 (2022).
- [17] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller, Equations of state calculations by fast computing machines, *J. Chem. Phys.* **21**, 1087 (1953).
- [18] V. F. Turchin, V. P. Kozlov, and M. S. Malkevich, The use of mathematical-statistics methods in the solution of incorrectly posed problems, *Sov. Phys. Usp.* **13**, 681 (1971).
- [19] D. L. Phillips, A technique for the numerical solution of certain integral equations of the first kind, *J. ACM* **9**, 84 (1962).
- [20] D. Frenkel, Speed-up of Monte Carlo simulations by sampling of rejected states, *Proc. Natl. Acad. Sci.* **101**, 17571 (2004).
- [21] A. Gelman, W. R. Gilks, and G. O. Roberts, Weak convergence and optimal scaling of random walk Metropolis algorithms, *Ann. Appl. Probab.* **7**, 110 (1997).
- [22] H. Haario, E. Saksman, and J. Tamminen, An adaptive Metropolis algorithm, *Bernoulli* **7**, 223 (2001).
- [23] C. Andrieu and J. Thoms, A tutorial on adaptive MCMC, *Stat. Comput.* **18**, 343 (2008).
- [24] J. Goodman and W. J., Ensemble samplers with affine invariance, *Comm. Appl. Math. Comp. Sci.* **5**, 65 (2010).
- [25] M. Vihola, Robust adaptive Metropolis algorithm with coerced acceptance rate, *Stat. Comput.* **22**, 997 (2012).
- [26] M. Maiberg, T. Hölscher, S. Zahedi-Azad, and R. Scheer, Theoretical study of time-resolved luminescence in semiconductors. III. Trap states in the band gap, *J. Appl. Phys.* **118**, 105701 (2015).
- [27] D. Kuciauskas, J. V. Li, A. Kanevce, H. Guthrey, M. Contreras, J. Pankow, P. Dippo, and K. Ramanathan, Charge-carrier dynamics in polycrystalline thin-film $\text{CuIn}_{1-x}\text{Ga}_x\text{Se}_2$ photovoltaic devices after pulsed laser excitation: Interface and space-charge region analysis, *J. Appl. Phys.* **117**, 185102 (2015).
- [28] E. Oksenberg, C. Fai, I. G. Scheblykin, E. Joselevich, E. L. Unger, T. Unold, C. Hages, and A. Merdasa, Deconvoluting energy transport mechanisms in metal halide perovskites using CsPbBr_3 nanowires as a model system, *Adv. Funct. Mater.* **31**, 2010704 (2021).
- [29] N. V. Smith, Classical generalization of the Drude formula for the optical conductivity, *Phys. Rev. B* **64**, 155106 (2001).
- [30] Y. Zhai, K. Wang, F. Zhang, C. Xiao, A. H. Rose, K. Zhu, and M. C. Beard, Individual electron and hole mobilities in lead-halide perovskites revealed by noncontact methods, *ACS Energy Lett.* **5**, 47 (2020).
- [31] R. Ulbricht, E. Hendry, J. Shan, T. F. Heinz, and M. Bonn, Carrier dynamics in semiconductors studied with time-resolved terahertz spectroscopy, *Rev. Mod. Phys.* **83**, 543 (2011).
- [32] M. M. Taheri, T. M. Truong, S. Li, W. N. Shafarman, B. E. McCandless, and J. B. Baxter, Distinguishing bulk and surface recombination in CdTe thin films and solar cells using time-resolved terahertz and photoluminescence spectroscopies, *J. Appl. Phys.* **130**, 163104 (2021).
- [33] E. Serpetzoglou, I. Konidakis, G. Kakavelakis, T. Maksudov, E. Kymakis, and E. Stratakis, Improved carrier transport in perovskite solar cells probed by femtosecond transient absorption spectroscopy, *ACS Appl. Mater. Interfaces* **9**, 43910 (2017).
- [34] M. N. Rosenbluth and A. W. Rosenbluth, Monte Carlo calculation of the average extension of molecular chains, *J. Chem. Phys.* **23**, 356 (1955).
- [35] W. K. Hastings, Monte Carlo sampling methods using Markov chains and their applications, *Biometrika* **57**, 97 (1970).
- [36] <https://doi.org/10.5281/zenodo.8178610>.
- [37] <https://github.com/HagesLab/MetroTRPL>.