

Bayesian Probabilistic Inference of Nonparametric Distance Distributions in DEER Spectroscopy

Sarah R. Sweger, Julian C. Cheung, Lukas Zha, Stephan Pribitzer, and Stefan Stoll*

Department of Chemistry, University of Washington, Seattle, WA 98195, United States

E-mail: stst@uw.edu

1 Abstract

2 Double Electron–Electron Resonance (DEER) 30 Double Electron–Electron Resonance (DEER)
3 spectroscopy measures distance distributions 31 spectroscopy is a pulse Electron Paramagnetic
4 between spin labels in proteins, yielding im- 32 Resonance (EPR) technique utilized for deter-
5 portant structural and energetic information 33 mining distances between electron spin centers
6 about conformational landscapes. Analysis of 34 on a nanometer scale,^{1,2} predominantly on pro-
7 an experimental DEER signal in terms of a dis- 35 teins. DEER resolves the full distribution of
8 tance distribution is a nontrivial task due to the 36 distances in an ensemble of proteins, making
9 ill-posed nature of the underlying mathemat- 37 it possible to directly quantify conformational
10 ical inversion problem. This work introduces 38 ensembles and the underlying conformational
11 a Bayesian probabilistic inference approach to 39 landscapes.^{3–5} DEER measures an oscillatory
12 analyze DEER data, assuming a nonparametric 40 time-domain signal that depends on the magni-
13 distance distribution with a Tikhonov smooth- 41 tude of the magnetic dipole–dipole interactions
14 ness prior. The method uses Markov Chain 42 between the spin centers, which in turn depends
15 Monte Carlo (MCMC) sampling with a compo- 43 on the inverse cube of the distance r . In the
16 sitional Gibbs sampler to determine a posterior 44 analysis, this signal is fitted with a model that
17 probability distribution over the entire param- 45 includes a distance distribution $P(r)$. Mathe-
18 eter space, including the distance distribution, 46 matically, this constitutes an ill-posed inversion
19 given an experimental dataset. This posterior 47 problem. Assessment of uncertainty in the fit-
20 contains all the information available from the 48 ted distance distribution is therefore challeng-
21 data, including a full quantification of the un- 49 ing, but is crucial for making sound conclusions
22 certainty about the model parameters. The cor- 50 about the conformational landscape of the pro-
23 responding uncertainty about the distance dis- 51 tein.
24 tribution is captured via an ensemble of poste- 52 Analysis approaches for obtaining a distance
25 rior predictive distributions. Several examples 53 distribution from a measured DEER signal
26 are presented to illustrate the method. Com- 54 range from analytical solutions⁶ to neural net-
27 pared to bootstrapping, it performs faster and 55 works.^{7,8} The least-squares fitting methods that
28 provides slightly larger uncertainty intervals. 56 have seen the widest practical application uti-
57 lize one of two models for $P(r)$: either a Gaus-
58 sian mixture model, or a non-parametric rep-
59 resentation combined with Tikhonov regular-

29 1 Introduction

30 Double Electron–Electron Resonance (DEER) 31 spectroscopy is a pulse Electron Paramagnetic
32 Resonance (EPR) technique utilized for deter-
33 mining distances between electron spin centers
34 on a nanometer scale,^{1,2} predominantly on pro-
35 teins. DEER resolves the full distribution of
36 distances in an ensemble of proteins, making
37 it possible to directly quantify conformational
38 ensembles and the underlying conformational
39 landscapes.^{3–5} DEER measures an oscillatory
40 time-domain signal that depends on the magni-
41 tude of the magnetic dipole–dipole interactions
42 between the spin centers, which in turn depends
43 on the inverse cube of the distance r . In the
44 analysis, this signal is fitted with a model that
45 includes a distance distribution $P(r)$. Mathe-
46 matically, this constitutes an ill-posed inversion
47 problem. Assessment of uncertainty in the fit-
48 ted distance distribution is therefore challeng-
49 ing, but is crucial for making sound conclusions
50 about the conformational landscape of the pro-
51 tein.

52 Analysis approaches for obtaining a distance
53 distribution from a measured DEER signal
54 range from analytical solutions⁶ to neural net-
55 works.^{7,8} The least-squares fitting methods that
56 have seen the widest practical application uti-
57 lize one of two models for $P(r)$: either a Gaus-
58 sian mixture model, or a non-parametric rep-
59 resentation combined with Tikhonov regular-

ization.^{9–16} Gaussian mixture models are parametric and represent the distribution as a linear combination of several Gaussian functions.¹¹¹ Non-parametric models represent $P(r)$ as a histogram over a discretized distance range, combined with Tikhonov regularization that includes a roughness penalty for the distribution into the fitting objective function. Both Gaussian mixture models and non-parametric models with Tikhonov regularization can be fit directly to the raw data in a single step.¹⁷ In both approaches, however, correctly quantifying and visualizing uncertainty is challenging.

For Gaussian mixture models, uncertainty analysis relies primarily on parameter confidence intervals, which are obtained from the covariance matrix or by explicitly exploring the sensitivity of the objective function on the parameter values.^{12–14,18} The parameter confidence intervals are then propagated to the distance domain to yield error bands on the distribution. This method assumes that the error surface is quadratic and that the parameters are unbounded, neither of which is generally true. Our previous work extended the available uncertainty analysis methods for Gaussian mixture models by implementing a Bayesian inference approach.¹⁶ The method models the raw DEER data directly and yields a full joint probability distribution over all model parameters, thereby fully quantifying their uncertainties. Similar Bayesian data analysis methods have been implemented previously for NMR and EPR.^{16,19–23}

For non-parametric distribution models with Tikhonov regularization, partial uncertainty analysis is commonly conducted by manually varying some parameters in the analysis (intermolecular background parameters, modulation depth, noise) and summarizing the sensitivity of the extracted distance distribution to these parameters into error bands around the fitted distribution.^{11,24} Another method to obtain confidence intervals for both approaches is bootstrapping. As implemented in DeerLab,¹⁷ it generates an ensemble of distributions by analyzing a large number of synthetically generated hypothetical signals based on the fitted model. In previous work, we introduced

a partial approach based on Bayesian inference to quantify the uncertainty in the distribution due to the noise in the signal.²² Unfortunately, this work required prior processing to remove the intermolecular background contribution and could not incorporate parameters beyond noise.

In this paper, we present a Bayesian inference approach for analyzing a DEER trace using a non-parametric distance distribution with Tikhonov smoothing. It extends our previous work²² and models the raw DEER data directly without prior background correction. It yields a full probability distribution for all model parameters, providing complete quantitative information about uncertainty and correlations for all parameters, without any implicit limiting assumptions. We also introduce distribution ensembles as a visual tool to effectively represent uncertainty about the distance distribution, including correlations which are neglected when using visualizations based on error bands.

The paper is structured as follows. Section 2 presents the model and section 3 outlines the inference methodology. Section 4 shows examples using synthetic and experimental data, including a comparison between parametric and nonparametric $P(r)$ models, an analysis of the dependence on the distance range, and a comparison of the quantified uncertainty with that obtained from bootstrapping. Finally, section 5 discusses the merits of this method in comparison to others.

2 Probabilistic DEER model

To model the DEER data, we start from the general noise-free continuous-time physical model^{17,25}

$$V_M(t) = V_0 \int_0^\infty K(t, r) P(r) dr \quad (1)$$

where $V_M(t)$ represents the DEER signal as a function of dipolar evolution time t , V_0 is an overall amplitude factor, $P(r)$ is the distribution of intramolecular inter-spin distances r normalized such that $\int_0^\infty P(r) dr = 1$, and $K(t, r)$ is the kernel function that provides the

152 DEER signal as a function of t and r . In this 187 or, equivalently,
 153 work, we use

$$K(t, r) = [(1 - \lambda) + \lambda K_0(t, r)] \cdot B(t) \quad (2)$$

154 with the modulation depth λ , the elementary 189 kernel function

$$K_0(t, r) = \int_0^1 \cos [(1 - 3 \cos^2 \theta) D r^{-3} t] d \cos \theta \quad (3)$$

156 and the constant $D = (\mu_0/4\pi)g^2\mu_B^2/\hbar$, with the 188 electron g factor g , the spin concentration c , 157 the Bohr magneton μ_B , the reduced Planck 190 constant \hbar , and the magnetic constant μ_0 . θ is the 159 angle between the applied magnetic field and 192 the inter-spin direction.

160 The first factor in Eq. (2) represents the 193 intramolecular contribution. The second factor, 194 $B(t)$, represents the intermolecular contribu- 195 tion, also called the background and sometimes 196 denoted as $V_{\text{inter}}(t)$. Here, we use an expo- 197 nential decay corresponding to a homogeneous 198 three-dimensional distribution of spins

$$B(t) = \exp(-k|t|) \quad (4)$$

169 with $k = (8\pi^2/9\sqrt{3})Dc\lambda$. More extended 207 background models that incorporate fractal di- 170 199 mensions or volume exclusions are occasionally 200 needed and can be incorporated easily,²⁶ al- 201 though they will increase the number of model 202 parameters. We will use the value of the back- 203 ground function at the end of the DEER time 204 trace

$$B_{\text{end}} = B(t_{\text{end}}) \quad (5)$$

177 as an alternate way to specify the background 205 decay rate, via $k = -\log(B_{\text{end}})/t_{\text{end}}$.

179 Experimentally, the DEER signal is mea- 206 sured at a set of linearly spaced discrete time 207 points $t = t_i$ and is therefore represented 212 as a n_t -element vector \mathbf{V} with elements V_i . 213 The measured values typically include normally 214 distributed noise ϵ_i with mean zero and t - 215 independent variance σ^2 .²² This is represented 216 as

$$V_i = V_{\text{M}}(t_i) + \epsilon_i \quad \text{with} \quad \epsilon_i \sim \text{normal}(0, \sigma^2) \quad (6)$$

$$V_i \sim \text{normal}(V_{\text{M}}(t_i), \tau^{-1}) \quad (7)$$

188 The tilde \sim indicates that the quantity on the 189 left is distributed according to the probability 190 distribution on the right. The precision τ is the 191 inverse of the variance, $\tau = 1/\sigma^2$.

192 While it is possible to use a closed-form ex- 193 pression for the angular integral in Eq. (3), the 194 integral over the distance distribution in Eq. (1) 195 can only be approximated numerically. We do 196 this by discretizing $P(r)$ at a set of equidistant 197 distances $r = r_j$, giving the n_r -element vector 198 \mathbf{P} . This gives

$$\mathbf{V} = \mathbf{V}_{\text{M}} + \boldsymbol{\epsilon} = \mathbf{K}\mathbf{P} + \boldsymbol{\epsilon} \quad (8)$$

199 with the kernel matrix \mathbf{K} with elements $K_{i,j} = 200 V_0 K(t_i, r_j) \Delta r$, where $\Delta r = r_{j+1} - r_j$ is the 201 increment in the r domain. $\boldsymbol{\epsilon}$ is the noise vector 202 with elements ϵ_i .

203 Extracting the distance distribution from 204 Eq. (8) using Tikhonov regularization is done 205 by minimizing an objective function that 206 includes a misfit term and a Tikhonov regular- 207 ization term,

$$\mathbf{P}_\alpha = \min_{\mathbf{P} \geq 0} (\|\mathbf{V} - \mathbf{K}\mathbf{P}\|^2 + \alpha^2 \|\mathbf{L}\mathbf{P}\|^2) \quad (9)$$

208 Here, α is the Tikhonov regularization parameter 209 and \mathbf{L} is an operator matrix, most com- 210 monly the $(n_r - 2) \times n_r$ second-order difference 211 matrix

$$\mathbf{L} = \begin{pmatrix} 1 & -2 & 1 & & & 0 \\ & 1 & -2 & 1 & & \\ & & \ddots & \ddots & \ddots & \\ 0 & & & 1 & -2 & 1 \end{pmatrix} \quad (10)$$

212 With \mathbf{L} as defined, the endpoints of the dis- 213 tance range are neglected in this penalty term. 214 They can be included by extending \mathbf{L} with an 215 additional first row with -2 and 1 as the first 216 elements and with an additional last row with 217 1 and -2 as the last two elements.

218 The model specified above depends on a set 219 of parameters: the distance distribution vec- 220 tor penalized for roughness \mathbf{P}_α , the modulation 221 depth λ , the end point B_{end} of the intermolec-

222 ular background decay, the noise precision τ ,
 223 and the overall amplitude V_0 . We indicate the
 224 parameter set as

$$\boldsymbol{\theta} = \{\mathbf{P}_\alpha, \lambda, B_{\text{end}}, \tau, V_0\} \quad (11)$$

225 Although the model is formulated here in terms
 226 of B_{end} and τ for more efficient numerical sam-
 227 pling, we will show the results in terms of k
 228 and σ to remain consistent with the standard
 229 practices in DEER data analysis.

230 The goal of analyzing \mathbf{V} given the chosen
 231 model is to determine the range of values for the
 232 model parameters that are consistent with the
 233 data. Now, since \mathbf{V} is an incomplete represen-
 234 tation of $V(t)$ due to time truncation, time dis-
 235cretization, amplitude discretization, and am-
 236 plitude noise, there will be uncertainty associ-
 237 ated with these parameters. It is important to
 238 quantify this uncertainty as well. Therefore, the
 239 goal of the analysis is to determine the full joint
 240 probability distribution of all parameters, given
 241 the data \mathbf{V} , symbolically written as

$$p(\boldsymbol{\theta}|\mathbf{V}) \quad (12)$$

242 This distribution is called the posterior distri-
 243 bution, or simply posterior. It is posterior in
 244 the sense that it represents the probability dis-
 245 tribution of the parameters *after* the data are
 246 taken into account. Once calculated, the poste-
 247 rior can be visualized or used to obtain statistics
 248 on the parameters, such as means and spreads.

249 Using Bayes' theorem, the posterior can be
 250 calculated via

$$p(\boldsymbol{\theta}|\mathbf{V}) \propto p(\mathbf{V}|\boldsymbol{\theta}) \cdot p(\boldsymbol{\theta}) \quad (13)$$

251 The first term on the right provides the prob-
 252 ability of the data given specific values for the
 253 model parameters. If seen as a function of the
 254 parameters, it is called the likelihood function.
 255 Based on Eqs. (8) & (9), it is

$$p(\mathbf{V}|\boldsymbol{\theta}) = \text{normal}(\mathbf{V}; \mathbf{K}\mathbf{P}, \tau^{-1}) \quad (289)
 \propto \tau^{nt/2} \exp\left(-\frac{1}{2}\tau\|\mathbf{V} - \mathbf{K}\mathbf{P}\|^2\right) \quad (14)$$

256 The second term on the right in Eq. (13) is the
 257 prior distribution or simply prior. It represents

258 the probability distribution of the parameters
 259 *prior* to taking the data into account, summa-
 260 rizing information about the parameters that
 261 is available independently of the given dataset.
 262 For example, we know without any data that
 263 P_i is nonnegative, that λ is between 0 and 1,
 264 that V_0 is around 1 (assuming the experimental
 265 trace is normalized to maximum 1), and that
 266 B_{end} is between 0 and 1. We take the prior
 267 as a product of independent distributions over
 268 individual parameters:

$$p(\boldsymbol{\theta}) = p(\mathbf{P}, \delta)p(\lambda)p(V_0)p(\tau)p(B_{\text{end}}) \quad (15)$$

269 with the smoothness hyperparameter δ that is
 270 related to the Tikhonov regularization parame-
 271 ter α by²²

$$\delta = \alpha^2\tau \quad (16)$$

272 We include δ as an additional parameter in $\boldsymbol{\theta}$.
 273 Note that we treat B_{end} (or k) as an indepen-
 274 dent parameter, even though it depends on λ
 275 (see Eq. 4). Alternatively, the concentration c
 276 could be used as a model parameter instead of
 277 k or B_{end} .

278 As prior for \mathbf{P} , we encode our knowledge
 279 that \mathbf{P} is element-wise non-negative, normal-
 280 ized, and expected to be smooth, i.e. P_i and P_j
 281 should be similar if the distances r_i and r_j are
 282 similar. For this, we write

$$p(\mathbf{P}, \delta) = p(\mathbf{P}|\delta)p(\delta)f(\mathbf{P}) \quad (17)$$

283 The function $f(\mathbf{P})$ in Eq. (17) is an indicator
 284 function that equals one if all elements of \mathbf{P} are
 285 non-negative and \mathbf{P} integrates to 1, and zero
 286 otherwise.

287 For the smoothness prior, we assume a normal
 288 distribution

$$p(\mathbf{P}|\delta) = \text{normal}(\mathbf{P}; \mathbf{0}, (\delta\mathbf{L}^T\mathbf{L})^{-1}) \quad (290)
 \propto \delta^{n/2} \exp\left(-\frac{1}{2}\delta\|\mathbf{L}\mathbf{P}\|^2\right) \quad (18)$$

289 where n is the number of non-zero elements in
 290 \mathbf{P} (to ensure proper normalization^{27,28}). This
 291 distribution assigns high prior probabilities to
 292 smooth distributions (where $\|\mathbf{L}\mathbf{P}\|^2$ is small)
 293 and low prior probabilities to rough distribu-
 294 tions (where $\|\mathbf{L}\mathbf{P}\|^2$ is large). This is motivated

physically by the flexibility of the spin labels and the biomacromolecule to which the pair of spin labels is attached. This corresponds to the penalty term in Eq. (9).

For the hyperprior for the regularization parameter δ , we select a gamma distribution²²

$$p(\delta) = \text{gamma}(\delta; a_\delta, b_\delta) \quad (19)$$

with $a_\delta = 1$ and $b_\delta = 10^{-6}$. This gives a very broad distribution function that decays exponentially with increasing δ .

For the priors for the other parameters, we use the same broad distributions as used in our earlier work on parametric \mathbf{P} models¹⁶

$$p(\lambda) = \text{beta}(\lambda; 1.3, 2) \quad (20)$$

$$p(B_{\text{end}}) = \text{beta}(B_{\text{end}}; 1.0, 1.5 \mu\text{s}) \quad (21)$$

$$p(V_0) = \text{bnd}(\text{normal}(V_0; 1, 0.2^2), 0) \quad (22)$$

$$p(\tau) = \text{gamma}(\tau; a_\tau, b_\tau) \quad (23)$$

with $a_\tau = 1$ and $b_\tau = 10^{-4}$. However, the priors do not necessarily need to follow these distributions. Instead, they should be chosen based on the known information about the system and setup. The same applies for the hyperparameter δ which will be reported as $\lg(\alpha)$ throughout using Eq. (16). The estimation of posterior probabilities from prior probabilities and observed data through Bayesian inference is visualized in Fig. 1.

With the above expressions, the posterior $p(\boldsymbol{\theta}|\mathbf{V})$ is fully defined. It has some structure that is important to recognize. (a) It is a gamma distribution in τ

$$p(\tau|\mathbf{V}, \boldsymbol{\theta}_{-\tau}) = \text{gamma}(\tau; \tilde{a}_\tau, \tilde{b}_\tau) \quad (24)$$

where $\boldsymbol{\theta}_{-\tau}$ indicates the set of all parameters except τ . The distribution parameters are $\tilde{a}_\tau = a_\tau + n_t/2$ and $\tilde{b}_\tau = b_\tau + \|\mathbf{V} - \mathbf{KP}\|^2/2$, where n_t is the number of elements in \mathbf{V} . (b) It is a gamma distribution in δ

$$p(\delta|\mathbf{V}, \boldsymbol{\theta}_{-\delta}) = \text{gamma}(\delta; \tilde{a}_\delta, \tilde{b}_\delta) \quad (25)$$

with $\tilde{a}_\delta = a_\delta + n/2$ and $\tilde{b}_\delta = b_\delta + \|\mathbf{LP}\|^2/2$. (c) It is a truncated multivariate normal distribu-

tion in \mathbf{P}

$$p(\mathbf{P}|\mathbf{V}, \boldsymbol{\theta}_{-\mathbf{P}}) = \text{normal}(\mathbf{P}; \bar{\mathbf{P}}, \boldsymbol{\Sigma}) f(\mathbf{P}) \quad (26)$$

with center $\bar{\mathbf{P}} = \tau \boldsymbol{\Sigma} \mathbf{K}^T \mathbf{V}$ and covariance matrix $\boldsymbol{\Sigma} = (\tau \mathbf{K}^T \mathbf{K} + \delta \mathbf{L}^T \mathbf{L})^{-1}$. We will make use of these structures for the sampling methodology in the next section.

3 Inference

The analytical form of the posterior distribution $p(\boldsymbol{\theta}|\mathbf{V})$ is intractable. In particular, it is not possible to evaluate integrals required to determine the mean, the variance, or marginalized distributions of individual parameters. Therefore, we resort to representing the distribution by a finite set of samples generated numerically, such that the density of samples in various regions in parameter space is proportional to the local probability density (see Fig. 2). These samples are then used to evaluate (approximately) the aforementioned integrals and to construct visualizations.

We use Markov Chain Monte Carlo (MCMC) sampling to generate a Markov chain of samples from the posterior,²⁹ where each sample i , containing $(\mathbf{P}_i, \tau_i, \delta_i, V_{0,i}, B_{\text{end},i}, \lambda_i)$, is generated from the previous sample $i-1$, containing $(\mathbf{P}_{i-1}, \tau_{i-1}, \delta_{i-1}, V_{0,i-1}, B_{\text{end},i-1}, \lambda_{i-1})$. A simplified example of the MCMC sampling process is shown in Fig. 2.

We investigated two separate MCMC sampling strategies. The first strategy involves a compositional Gibbs sampling approach that utilizes independent draws for the three parameters that have simple analytical conditional posterior distributions (τ , δ , and \mathbf{P}) and Hamiltonian Monte Carlo (HMC) draws for the remaining three (V_0 , B_{end} , λ). The second one utilizes HMC sampling for all model parameters.

The compositional Gibbs sampling approach is based on our previous work.²² After choosing a starting point, it proceeds iteratively as follows:

- (1) Generate the i -th random sample of the precision τ from its full conditional posterior

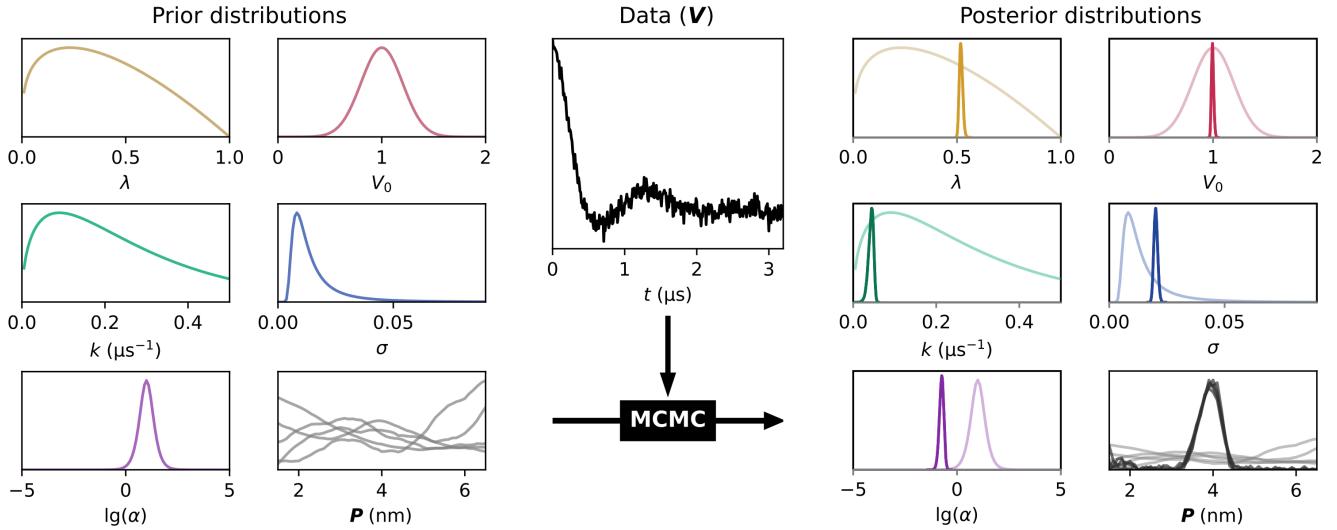


Figure 1: Visualization of Bayesian inference. Prior distributions for model parameters (modulation depth (λ), echo amplitude (V_0), background decay constant ($k = -\log(B_{\text{end}})/t_{\text{end}}$), noise ($\sigma = 1/\sqrt{\tau}$), Tikhonov smoothness parameter ($\alpha = \sqrt{\delta/\tau}$), and the distance distribution (\mathbf{P}) are combined with the data \mathbf{V} to yield a posterior probability distribution, with the marginalized posterior distributions for individual parameters shown. An ensemble of 5 distance distributions are shown for the prior and posterior of \mathbf{P} .

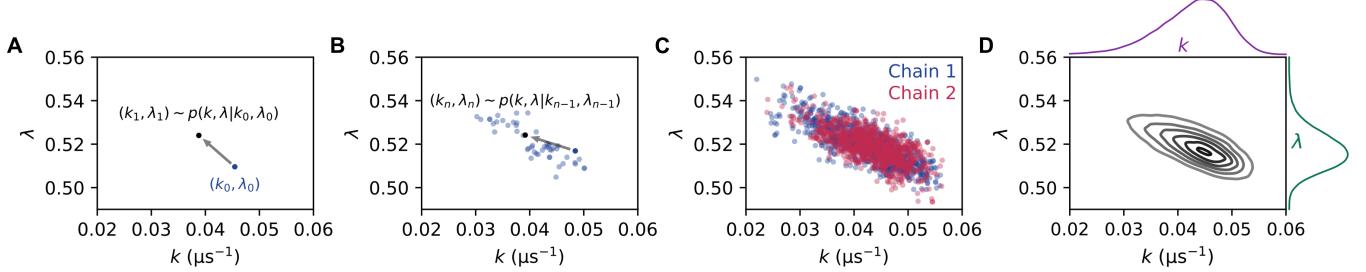


Figure 2: Principle of MCMC sampling of the posterior distribution. A. A random starting value is chosen from the joint posterior of k and λ (blue). The next random sample is chosen from the probability distribution of k and λ given the starting value (black), and thus every step is dependent upon the previous step. B. This process continues, where each new step is a random draw from a probability distribution for k and λ given the position of the last step. C. Regions of higher probability are sampled more frequently until the chain converges. Multiple chains are run, indicated by color, to obtain inter- and intra-chain convergence. D. The result of the sampling is a representation of the full posterior distribution of the parameters. Shown here are the 1D and 2D marginal posteriors for k and λ .

371 distribution Eq. (24) using the values of all 420
372 other parameters from sample $i - 1$. 421

373 (2) Generate the i -th random sample of the 422
374 regularization parameter δ from its posterior 423
375 distribution Eq. (25) using the new value for τ 424
376 and the previous values of all other parameters. 425

377 (3) Generate the i -th random sample of \mathbf{P} 424
378 from its full conditional posterior distribution 425
379 Eq. (26) with the new values for τ and δ and the 426
380 previous values for all other parameters from 427
381 sample $i - 1$. To generate a random sample from 428
382 this distribution, we use the fast non-negative 429
383 least-squares (FNNLS) algorithm by Bro and 430
384 De Jong³⁰ to enforce the non-negativity and 431
385 normalization constraints imposed by $f(\mathbf{P})$. 432
386 Other algorithms for generating samples from a 433
387 truncated multivariate normal distribution are 434
388 available in the literature.^{31,32} 434

389 For the remaining three parameters, since 435
390 their posteriors are not of a form for which 436
391 independent sampling is possible, we use an 437
392 HMC algorithm known as the no-U-turn sam- 438
393 pler (NUTS)³³ to simultaneously generate the 439
394 next samples of these parameters: 440

395 (4) Generate the i -th random sample of the 441
396 remaining parameters (V_0 , B_{end} , and λ) with 442
397 the NUTS sampler, using the new values of \mathbf{P} , 443
398 τ and δ and the previous values of V_0 , B_{end} and 444
399 λ . 445

400 HMC methods take the negative logarithm 446
401 of the posterior distribution to draw from as 447
402 a potential-energy landscape and the parame- 448
403 ters as position variables of fictitious particles. 449
404 Samples are then generated by simulating par- 450
405 ticle trajectories on this landscape with classi- 451
406 cal Hamiltonian dynamics using momenta that 452
407 are drawn from a multivariate normal distribu- 453
408 tion.^{34,35} NUTS auto-tunes the step size and the 454
409 number of steps used in the integration of the 455
410 Hamiltonian dynamics. 456

411 The second sampling approach we investi- 457
412 gated samples all six parameters simultane- 458
413 ously with the NUTS sampler. However, since 459
414 Eq. (18) would yield negative values for \mathbf{P} 460
415 when sampled with NUTS, we represent \mathbf{P} as a 461
416 uniform Dirichlet distribution $\text{Dir}(\mathbf{P}; \mathbf{1})$ to en- 462
417 code the non-negativity and normalization con- 463
418 straints from $f(\mathbf{P})$ in Eq. (17). To accommo- 464
419 date the smoothness prior in this implementa- 465

tion, we add the $p(\mathbf{P}|\delta)$ term from Eq. (18) di- 420
rectly as an additional term to the potential- 421
energy function of the NUTS sampler. Thus, 422
Eq. (17) becomes:

$$p(\mathbf{P}, \delta) \propto p(\mathbf{P}|\delta)p(\delta)\text{Dir}(\mathbf{P}; \mathbf{1}) \quad (27)$$

We continue to use Eqs. (19) to (23) as our priors for the other parameters.

We implemented both sampling algorithms in the Python package PyMC 5.10.4,³⁶ which uses autodifferentiation for the calculation of the gradient necessary for calculating the Hamiltonian trajectories in NUTS sampling. For each analysis, 4 chains containing several thousand to several tens of thousands of samples are run. These chains are then assessed for convergence using the rank-normalized split \hat{R} statistic, which compares intra- and inter-chain variances.³⁷⁻³⁹ Values of \hat{R} very close to 1 indicate that the chains are stationary and similar, such as in the example in Fig. 2C. Chains in this work are considered converged when $\hat{R} < 1.05$. Although convergence is essential to proper analysis, some runs take impractically long to converge. When this occurs, we exclude chains one by one and observe the effect on \hat{R} , and then remove the chain that leads to the largest reduction in \hat{R} when excluded. In practice, we did not need to remove more than one chain to attain convergence, but this process can be repeated. All code used for modeling and sampling in this paper was run within our Python package dive, which can be accessed at <https://github.com/StollLab/dive>.

After convergence, the pooled samples from all chains represent the full N -dimensional posterior $p(\boldsymbol{\theta}|\mathbf{V})$. Due to its large dimensionality, it is not possible to visualize it directly. Instead, we examine each parameter individually using its marginalized posterior, which is obtained by integrating the full posterior over all other parameters. This integral is approximated by generating a histogram of the parameter values from all samples, smoothed with a Gaussian with a line width of 1/5 of the standard deviation of the parameter values. This results in a one-dimensional distribution that can easily be plotted. On the right of Fig. 1, the marginal-

ized posteriors are shown in color, together with the priors in less saturated color. In this case, the posteriors are much narrower than the corresponding priors. The spread of the posterior distribution is a quantitative measure of inferential uncertainty, and its narrowing compared to the prior is a direct measure of information content of the data.

However, marginalization discards all information about correlation between parameters. It is also possible to display and examine two-dimensional marginalized posteriors between pairs of parameters, which is particularly helpful for assessing issues of convergence often caused by highly correlated parameters. Examples of 1D and 2D marginalized posteriors are shown in Fig. 2D.

We additionally visualize the results of the Bayesian inference using posterior predictive samples for $V(t)$ and $P(r)$. For this, we randomly pick a small set of samples (about 50-100) of the parameters, including the distance distribution, from the pooled MCMC samples. An ensemble of noise-free time-domain signals and background decays is then generated from the sampled parameters. Plotting these posterior predictive samples of the time-domain trace and the distance distribution allows for a visual assessment of fit quality and of uncertainty in the inferred distance distribution.

When visualized, we found that the two MCMC methods we investigated did have some differences. Specifically, when using the NUTS sampler for all parameters, the sampled distance distributions differ from those of the compositional Gibbs–NUTS sampler: whereas the compositional sampler generates \mathbf{P} vectors with a significant number of points equal to exactly 0, the NUTS sampler never generates \mathbf{P} with points equalling 0.

This difference arises from the use of the FNNLS algorithm to generate non-negative \mathbf{P} draws in the compositional sampling approach. The FNNLS algorithm initially sets all points in \mathbf{P} to zero in a non-negativity constrained “active set” and improves the fit by iteratively moving points to an unconstrained “passive set” until the fit can no longer be improved.³⁰ At this point, \mathbf{P} consists of some positive points

and some zero points, meaning that the probability of points in \mathbf{P} being equal to 0 is significant. This is similar to a spike-and-slab prior, a common distribution in Bayesian inference involving a discontinuity at 0 in an otherwise smooth distribution to increase the probability of 0.²⁸ Spike-and-slab priors are examples of priors that encourage sparse distributions; the continuous Laplace, double Pareto, and horseshoe priors are other examples of sparsity-inducing priors.⁴⁰ Thus, the FNNLS algorithm effectively adds an implicit sparsity-inducing term to Eq. (17) that encourages points in \mathbf{P} to be 0, representing our knowledge (or assumption) that there are many distances that the spin label pair does not populate. Although a similar effect could be achieved in the NUTS sampler by using an explicit sparsity-inducing prior for \mathbf{P} in Eq. (27), few multidimensional, non-negative, constant-sum, sparsity-inducing priors are available in the literature.

Furthermore, in our implementation, the NUTS sampler runs more slowly than the Gibbs–NUTS sampler due to the larger number of parameters that are included in the calculation of the potential-energy landscape. We also found that the additional complexity of this landscape leads to a greater frequency of undesirable divergences, which occur when the NUTS sampler, which takes discrete steps, encounters regions that are too steep to sample accurately. For these reasons, we choose to use the compositional Gibbs–NUTS sampling method for the rest of this discussion, taking note of the implicit sparsity bias. However, the NUTS sampling approach remains as an alternative.

4 Results & Discussion

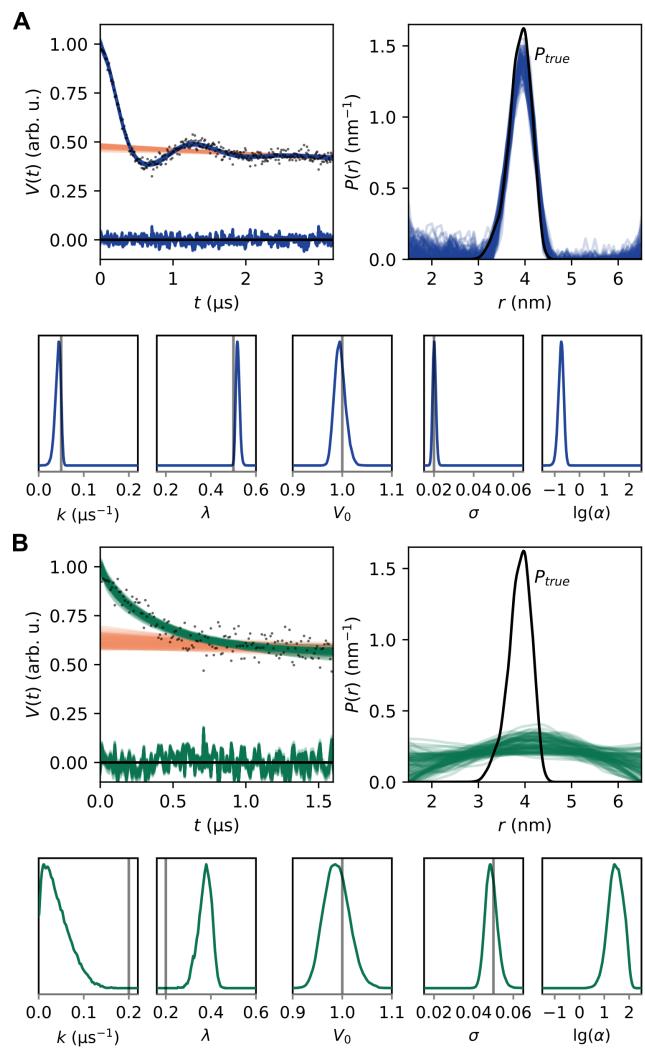
4.1 Basic illustrations

We first illustrate the probabilistic analysis method on synthetic data, using the large simulated T4 lysozyme (T4L) test data set published by Edwards et al.^{15,41} The distributions in this test data set were generated computationally from an in silico spin-labeled crystal

560 structure of T4L. Distribution 3992 from the
 561 test data set is taken as ground truth and two
 562 DEER traces of differing quality were generated
 563 and then analyzed with the Bayesian inference
 564 method using the aforementioned nonparamet-
 565 ric model. The traces and the analysis results
 566 are shown in Fig. 3. The trace in panel A has a
 567 large modulation depth, a slow background de-
 568 cay rate, a long trace length, and a small noise
 569 level. Comparatively, the trace in panel B has
 570 less ideal values for all of these parameters, in
 571 particular a shorter trace length and a higher
 572 noise level. This provides a challenging case
 573 with higher inferential uncertainty.

574 For both cases, the marginalized posteriors
 575 for all scalar parameters are shown (Fig. 3A &
 576 B, bottom). The gray lines indicate the ground-
 577 truth values used in generating the trace. There
 578 is no line shown for α as this is a non-physical
 579 parameter introduced in the analysis. For the
 580 longer and less noisy trace, V_0 and σ are recov-
 581 ered accurately and with little uncertainty, as
 582 indicated by the narrow posterior distributions
 583 with modes close to the ground-truth values.
 584 For k and λ , the posterior modes align less with
 585 the ground-truth values, and the spread in the
 586 case of k is a bit larger. The cause for this is
 587 discussed in more depth below.

588 The parameter posteriors most directly show
 589 the outcome of the Bayesian analysis and are
 590 useful for identifying the impact of individual
 591 parameters on the overall fit. However, the
 592 time-domain fit and the distance distribution
 593 are the results of primary interest. These quan-
 594 tities are shown via ensembles of posterior pre-
 595 dictive samples (Fig. 3A & B, top), as described
 596 in Section 3. As in our previous work, this
 597 visualization for uncertainty is preferred as it
 598 does not emphasize any particular distribution
 599 and more completely encompasses the range
 600 of fits compatible with the data. The time-
 601 domain plot show that the fit is excellent, and
 602 the distance-domain plots shows that there is
 603 little scatter in center, width and shape among
 604 the $P(r)$ distributions. The ensemble also over-
 605 laps well with the ground truth, indicating that
 606 the Bayesian analysis recovers the distribution
 607 from the data with little uncertainty. There
 608 is somewhat elevated uncertainty at short dis-



596 Figure 3: Validation of the Bayesian infer-
 597 ence method using synthetic data generated
 598 from distribution 3992 from the Edwards test
 599 set.^{15,41} The time trace in panel A was gen-
 600 erated using $\lambda = 0.5, k = 0.05 \mu\text{s}^{-1}, t_{\text{max}} =$
 $601 3.2 \mu\text{s}, \sigma = 0.02$, and the trace in panel B was
 602 generated using $\lambda = 0.2, k = 0.2 \mu\text{s}^{-1}, t_{\text{max}} =$
 $603 1.6 \mu\text{s}, \sigma = 0.05$. The MCMC simulation of
 604 each was run with 4 chains and 20,000 samples
 605 per chain. In each panel, the top left plot shows
 606 the time domain data (gray dots), ensembles of
 607 fitted signals \mathbf{V} (blue/green) and backgrounds
 608 $(1 - \lambda)\mathbf{B}$ (orange), and the residuals. The top
 609 right plot shows an ensemble of distance distri-
 610 butions drawn from the posterior distribution
 611 (blue/green) and the ground-truth distance dis-
 612 tribution (black). The bottom plots show the
 613 marginalized parameter posteriors. The gray
 614 lines indicate the true parameter values.

tances and at the long-distance edge. Analysis of the posterior modes for the shorter, noisier dataset (Fig. 3B) shows that for V_0 and σ the modes are reasonably well identified, but the spreads are wider than for the first dataset. For k and λ , the method is unable to recover the ground-truth values accurately. The difficulty of recovering these parameters is indicative of a problem inherent to the underlying physical model—parameter non-identifiability.⁴² The signals from the intermolecular background and from intramolecular long distances are very similar. Given short and noisy data, there is insufficient information for identifying and separating the two contributions, resulting in skewed and broad posterior distributions for k and λ . In addition, the inferred distance distribution is very broad and uncertain, preventing specific structural conclusions. This broadness indicates that it is dominated by the prior, and that the data did not provide significant additional information about $P(r)$.

Both datasets in Fig. 3 show posterior distributions of α . In least-squares fitting approaches, a single value of α is selected ad hoc or based on one of a series of criteria (L-curve, Akaike information criterion, etc.). In the inference approach presented here, the prior together with the data result in a distribution of likely α value, without the need to pick a particular value or criterion.

4.2 Model comparison

The rest of the examples presented utilize the DEER data recently published as part of a benchmark test and guidelines paper for DEER.⁴³ Four constructs of the *Yersinia* outer protein O (YopO) from *Yersinia enterocolitica* without its membrane anchor were measured by seven different labs. We use the data from lab B. YopO contains an α -helix that is 43 amino acids long, allowing for three site pairs to be chosen that encompass the short- (S585R1/Q603R1), mid- (V599R1/N624R1), and long- (Y588R1/N624R1) range distances accessible by DEER. A fourth site pair was chosen to include a spin label on a flexible

loop, giving rise to a very broad distribution (S353R1/Q635R1).

The results of using the Bayesian inference approach with both the nonparametric model for $P(r)$ described in this paper and the parametric multi-Gauss model from our previous work¹⁶ are shown in Fig. 4. For each panel, the darker colored, top ensembles are the results of using the nonparametric model and the lighter colored, bottom ensembles are the parametric model results with number of Gaussians indicated. For all four samples, convergence is achieved for both models and the distance distributions show good agreement to those previously published.⁴³ The uncertainty is relatively low given the small scatter of the posterior predictive distributions. Uncertainty increases at long distances, but does not affect interpretation of the primary features. The parametric results for all site pairs show good agreement with the nonparametric results. The presence of peaks with high uncertainty in the parametric models in Fig. 4A, B & C is indicative of limitations of the parametric approach, namely, imposing an underlying shape to the distance distribution. The multi-Gauss parametric model has more difficulty recovering distributions that have multiple, overlapping peaks of similar width or intensity. Distributions of this nature show larger correlations between distribution parameters, making exploring the parameter space significantly less efficient and convergence more difficult to achieve. This was shown previously when analyzing synthetic data.¹⁶

4.3 Dependence on distance axis

A non-parametric $P(r)$ is not entirely free of parameters—it depends on the fixed parameters that define the distance axis. These are the minimum distance r_{\min} , the maximum distance, r_{\max} , and, for a linear axis, the resolution Δr . This forces the distribution to be zero outside the distance range, and imposes a fixed resolution within the range. Figure 5 shows a series of results for the mid-range YopO dataset with different r_{\max} at a constant resolution Δr . Two effects of r_{\max} can be discerned. First, as the r_{\max}

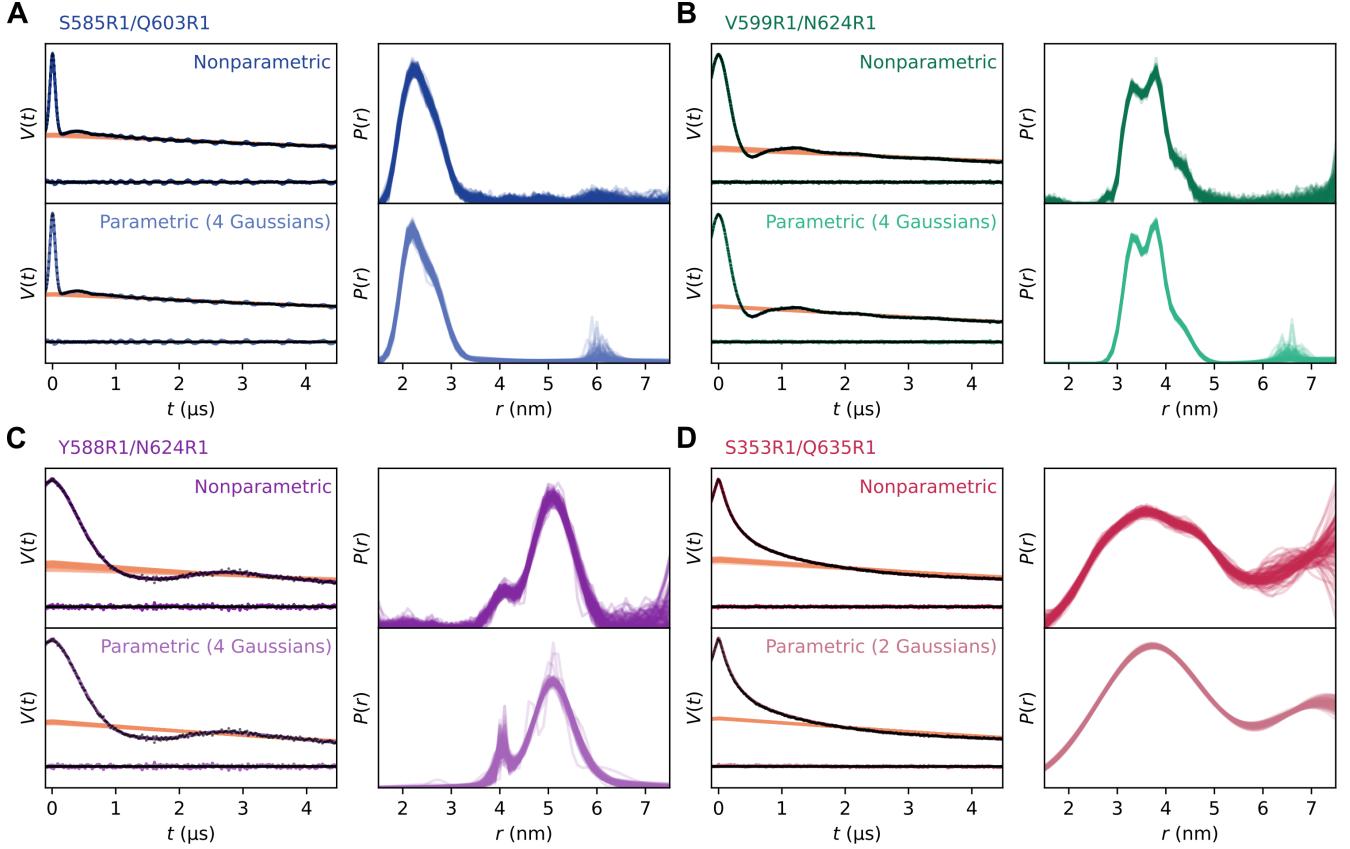


Figure 4: Bayesian inference using a nonparametric model with Tikhonov regularization and a multi-Gauss parametric model on DEER data obtained for four constructs of YopO.^{43,44} Most MCMC data shown were run with 4 chains of 20,000 samples per chain from which 100 samples were randomly selected. The nonparametric models for panels B, C, and D were run with 4 chains of 100,000 samples to achieve convergence according to $\hat{R} < 1.05$. For the nonparametric model for panel B and the parametric models for panels A, B, and D, one chain was dropped due to non-convergence. The raw experimental data are shown in gray overlaid with the time-domain full and intermolecular fits for the nonparametric model (top, dark) and the parametric multi-Gauss model (bottom, light). Beneath the time-domain data are the residuals from the posterior predictive ensemble. The distance distributions for each are shown according to the same color scheme.

703 is decreased, the uncertainty in the region near r_{\max} shrinks. Second, it becomes less arduous to 704 obtain converged chains (see Fig. 5 caption) and 705 the resulting posteriors indicate more efficient 706 sampling and certainty around the parameters. 707 When run for 20,000 samples per chain, the 708 MCMC sampler was only able to converge for 709 the model with an r_{\max} of 6.5 nm. The model 710 with an r_{\max} of 8.5 nm (which is past heuris- 711 tic values often used for r_{\max} , $\sqrt[3]{Dt_{\text{end}}/2\pi n_{\min}}$, 712 where n_{\min} is the number of required dipolar os- 713 cillation periods ($n_{\min} > 1/2$)),⁴³ did not attain 714 convergence, as evidenced by \hat{R} values much 715 greater than 1.05 even when run for 100,000 716 samples per chain. 717

718 Both effects are a consequence of the fact that 719 the models with long r_{\max} are over-specified— 720 there is not enough information in the data to 721 clearly distinguish between mass in this region 722 of r and the intermolecular background. The 723 marginal distributions of the background 724 parameter k and the modulation depth λ broaden 725 substantially as r_{\max} is increased. These ef- 726 fects are also noticeable in the accompany- 727 ing increased spread of background fits. Note 728 that these effects are not a peculiarity of the 729 MCMC method, but rather an intrinsic param- 730 eter non-identifiability issue of the model given 731 the data. Any analysis method that uses this 732 model will encounter similar difficulties in iden- 733 tifying unique parameter values. Based on the 734 posterior distributions, the Bayesian approach 735 provides a direct way for diagnosing these situ- 736 ations. 737

738 Several approaches can be considered to pre- 739 vent model overspecification: (a) restrict r_{\max} 740 to shorter values, (b) include additional infor- 741 mation that $P(r)$ is close to zero at long dis- 742 tances, (c) include additional information that 743 $P(r)$ is compact,⁴² or (d) use a less flexible 744 model with stronger assumptions about the 745 r distribution, such as a multi-Gauss model 746 or a model with significantly fewer distance 747 points.⁴⁵ 748

4.4 Comparison with bootstrapping

The Bayesian analysis presented here provides full quantification of uncertainty for all model parameters, including correlations. Another, although conceptually different, approach is bootstrapping, a Monte Carlo resampling method. In the bootstrapping variant implemented in DeerLab,¹⁷ synthetic data traces are generated by adding different noise realizations to a fitted signal obtained by least-squares fitting. These new traces are then analyzed according to the same procedure as the original experimental data. This results in a number of fitted parameter sets and distance distributions. The distribution of fitted values compares in nature to the posteriors output from Bayesian inference and can be randomly sampled from to produce ensemble plots representing the spread of uncertainty around a particular set of data.

In Fig. 6 we show the Bayesian analysis using a nonparametric model of the YopO data from above and compare it to the data analyzed with bootstrapping. Compared to our previous figures, the distance axis resolution was doubled to 0.05 nm and r_{\max} was set to 6.5 nm. For each panel, the darker colored, top ensembles are the results of using Bayesian inference and the lighter colored, bottom ensembles are the bootstrapping results. For bootstrapping, an initial fit was achieved using a regularization parameter selected by the Bayesian information criterion.¹⁵ This value of α was then frozen for the bootstrap analysis in which 1000 bootstrapped samples were taken, i.e., 1000 new signal traces were generated and fit. A set of 100 parameter vectors and distributions were randomly drawn from the 1000 samples and plotted alongside the Bayesian inference ensembles.

For all the site pairs, the ensembles of distance distributions are very similar between Bayesian analysis and bootstrapping (see Fig. 6). However, the bootstrap ensembles generally have less scatter, since all synthetic signals are generated from the same initial fit, leading to less exploration of the combined parameter space compared to the Bayesian approach.

A crucial difference between Bayesian infer-

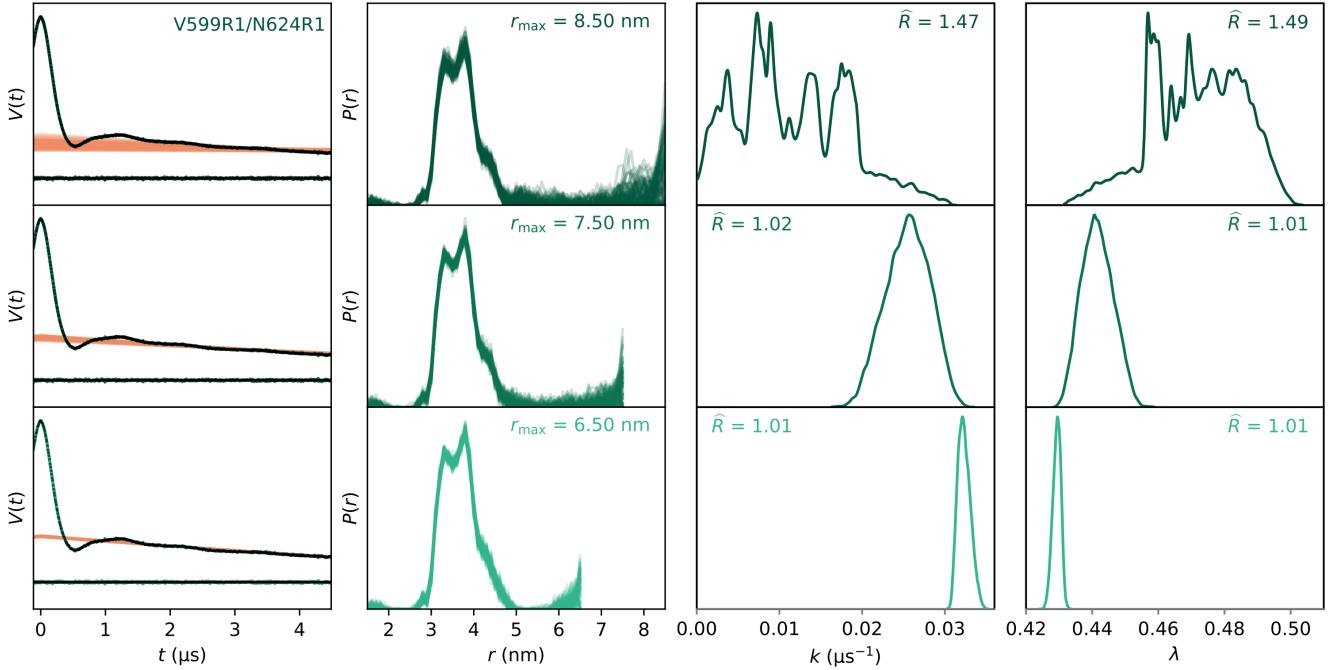


Figure 5: Effect of upper distance limit r_{\max} on quality of fit and $P(r)$ uncertainty. Data from the mid-range site pair (YopO V599R1/N624R1) was evaluated for three values of r_{\max} and a constant resolution (Δr) of 0.1 nm, using MCMC with 4 chains and 20,000 samples (for $r_{\max} = 6.5$ nm) or 100,000 samples ($r_{\max} = 7.5$ nm and 8.5 nm). One chain was then dropped from the 7.5 nm model to attain convergence. The 8.5 nm model did not converge regardless of how many chains were dropped, so all 4 chains were kept. From these runs, 100 distance distributions are randomly sampled and plotted. The raw experimental data are shown in gray overlaid with the associated time-domain fits. The saturation of the color decreases with r_{\max} . Left: Time-domain fits and resulting distance distributions. Right: The posterior distributions for the background decay rate constant k and the modulation depth λ with decreasing r_{\max} .

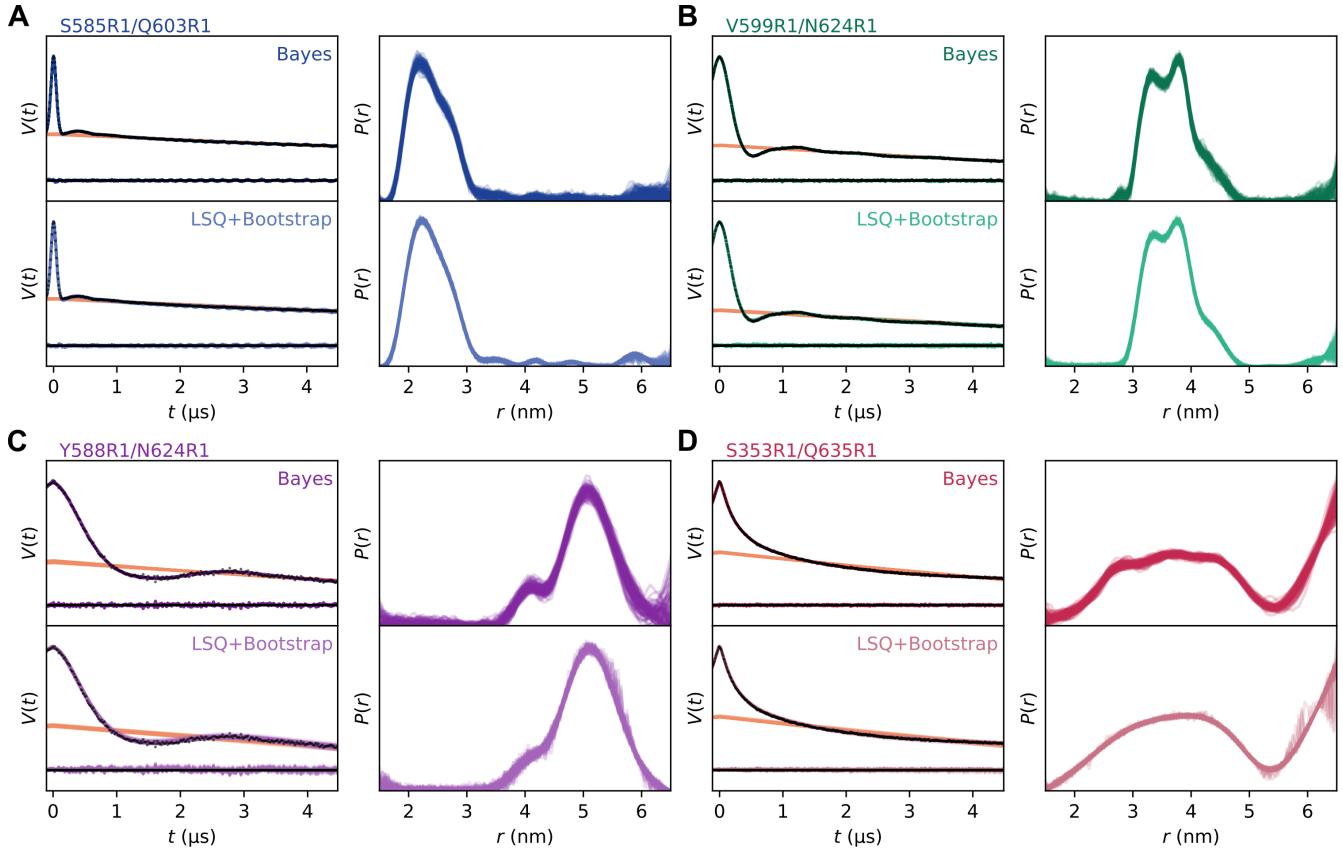


Figure 6: Comparison of experimental fits and visualization of uncertainty between Bayesian inference and bootstrapping. The MCMC data shown were run with 4 chains of 100,000 samples per chain. For both the Bayesian MCMC data and the bootstrapping fits, 100 samples were randomly selected to plot. The raw experimental data are shown in gray overlaid with the time-domain full and intermolecular fits for the Bayesian analysis (top, dark) and bootstrapping (bottom, light). The distance distributions for each are shown according to the same color scheme.

ence and bootstrapping is that in the latter method, the Tikhonov smoothing parameter α is fixed, whereas it is a floating (hyper)parameter in the former. Therefore, bootstrapping does not incorporate the uncertainty due to α . Yet, the value and uncertainty of α are crucial components for assessing whether over- or under-fitting is occurring and provide insight into the overall shape of the resulting distance distribution. With Bayesian inference, the uncertainty of α , along with all other model parameters, is a direct output of the analysis. This is illustrated in Fig. 7, which plots the marginalized posterior of $\lg(\alpha)$ in the Bayesian analysis against the fixed value of $\lg(\alpha)$ used in the bootstrap analysis (vertical line).

In terms of computational cost, the time taken to run 1,000 bootstrapped samples for the bootstrap analysis was similar to the time taken to run approximately 100,000 MCMC draws for the Bayesian analysis (a couple hours on a typical laptop computer). Thus, the Bayesian approach yields a more complete uncertainty analysis at roughly the same computational cost. Note, however, that computational performance depends on implementation, and that there is a wide range of possible bootstrapping approaches beyond the one implemented in DeerLab.

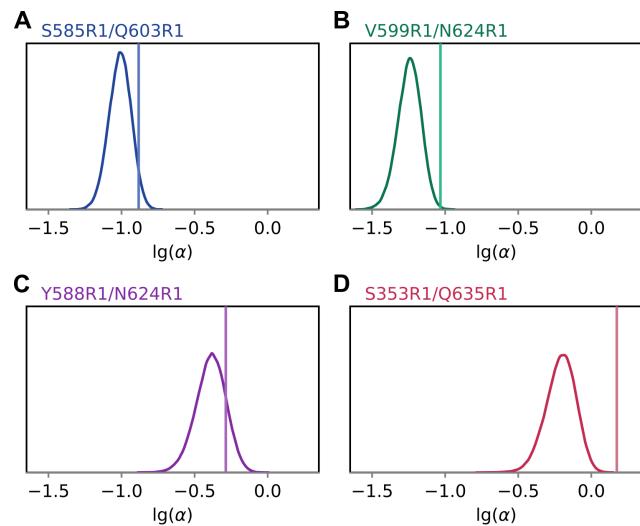


Figure 7: Uncertainty assessment for the Tikhonov smoothness parameter α . The posterior distributions for $\lg(\alpha)$ from the Bayesian analysis are shown, maintaining the same panel order and coloring as previous figures for the site pairs. The single values for the $\lg(\alpha)$ parameter from the bootstrap fits are shown as a vertical lines.

Gaussian model, could be advantageous, at the expense of biasing the analysis.

The complete uncertainty quantification provided by Bayesian inference has advantages over other forms of uncertainty quantification. Compared to confidence intervals based on maximum-likelihood estimation and the Laplace approximation (which assumes the likelihood function is Gaussian around its maximum), it is more complete and accurate and can capture asymmetric uncertainty as well as broad areas in parameter space that lead to equally good model fits. Therefore, it can help spot identifiability problems with the model parameters. Compared to bootstrapping, which can be used to determine more comprehensive confidence intervals within the maximum-likelihood estimation approach, Bayesian inference does not synthetically generate new datasets. The only dataset that is used in Bayesian inference is the given experimental dataset. From a principled standpoint, this is indeed the only dataset that should matter. Also, the Bayesian approach saves significant computation time on a per-sample basis. In

841 5 Conclusions

842 The Bayesian method outlined in this work 850 presents a rigorous, complete and conceptually 851 simple inference approach for analyzing DEER 852 data. Given the experimental data and the 853 choice of a particular physical model, it deter- 854 mines the joint probability distribution of all 855 model parameters. This provides, completely, 856 any information that can be gleaned from the 857 data under the assumption of the chosen model 858 and provided prior information. If the informa- 859 tion content of the data is low (noisy and trun- 860 cated trace), then this method captures the re- 861 sulting significant uncertainty about the model 862 parameters, particularly if a Tikhonov model 863 is used. Introducing additional constraints to 864 the r distribution, such as by using a multi- 865

866 general, the logically rigorous foundation of the 908
867 Bayesian inference approach makes it possible 909
868 to use it as a reference method for analyzing 910
869 dipolar EPR data.

870 Although we presented the method using a 911
 871 Tikhonov model for the intramolecular dis- 912
 872 tance distribution combined with a simple ex- 913
 873 ponential decay for the intermolecular contri- 914
 874 bution, this method is very general and can 915
 875 handle more complex models. Extensions to 916
 876 use more sophisticated intermolecular models, 917
 877 to include multiple dipolar pathways, and to 918
 878 handle multiple datasets simultaneously are 919
 879 conceptually straightforward, although compu- 920
 880 tationally more demanding. Therefore, the 921
 881 Bayesian frameworks provide a robust and flex- 922
 882 ible tool for DEER data analysis, providing 923
 883 complete quantitative information about uncer- 924
 884 tainty.

885 Acknowledgements

886 This work was supported by the National
887 Science Foundation with grant CHE-1452967 928
888 (S.S.) and by the National Institutes of Health 929
889 with grants R01 GM125753 and R01 GM127325 930
890 (S.S.). We thank Jonathan Bardsley (Univer- 931
891 sity of Montana) for helpful discussions. 932

892 References

general, the logically rigorous foundation of the Bayesian inference approach makes it possible to use it as a reference method for analyzing dipolar EPR data.

Although we presented the method using a Tikhonov model for the intramolecular distance distribution combined with a simple exponential decay for the intermolecular contribution, this method is very general and can handle more complex models. Extensions to use more sophisticated intermolecular models, to include multiple dipolar pathways, and to handle multiple datasets simultaneously are conceptually straightforward, although computationally more demanding. Therefore, the Bayesian frameworks provide a robust and flexible tool for DEER data analysis, providing complete quantitative information about uncertainty.

Acknowledgements

This work was supported by the National Science Foundation with grant CHE-1452967 (S.S.) and by the National Institutes of Health with grants R01 GM125753 and R01 GM127325 (S.S.). We thank Jonathan Bardsley (University of Montana) for helpful discussions.

References

- (1) Milov, A. D.; Salikhov, K. M.; Shchirov, M. D. Application of the Double Resonance Method to Electron Spin Echo in a Study of the Spatial Distribution of Paramagnetic Centers in Solids. *Sov. Phys. Solid State* **1981**, *23*, 565–569.
- (2) Milov, A. D.; Tsvetkov, Y. D. Double Electron–Electron Resonance in Electron Spin Echo: Conformations of Spin-Labeled Poly-4-Vinylpyridine in Glassy Solutions. *Appl. Magn. Reson.* **1997**, *12*, 495–504.
- (3) Pannier, M.; Veit, S.; Godt, A.; Jeschke, G.; Spiess, H. W. Dead-Time Free Measurement of Dipole–Dipole
- (4) Jeschke, G.; Polyhach, Y. Distance Measurements on Spin-Labelled Biomacromolecules by Pulsed Electron Paramagnetic Resonance. *Phys. Chem. Chem. Phys.* **2007**, *9*, 1895–1910.
- (5) Jeschke, G. DEER Distance Measurements on Proteins. *Annu. Rev. Phys. Chem.* **2012**, *63*, 419–446.
- (6) Matveeva, A. G.; Nekrasov, V. M.; Maryasov, A. G. Analytical Solution of the PELDOR Inverse Problem Using the Integral Mellin Transform. *Phys. Chem. Chem. Phys.* **2018**, *19*, 32381–32388.
- (7) Worswick, S. G.; Spencer, J. A.; Jeschke, G.; Kuprov, I. Deep Neural Network Processing of DEER Data. *Sci. Adv.* **2018**, *4*, eaat5218.
- (8) Keeley, J.; Choudhury, T.; Galazzo, L.; Bordignon, E.; Feintuch, A.; Goldfarb, D.; Russell, H.; Taylor, M. J.; Lovett, J. E.; Eggeling, A. et al. Neural networks in pulsed dipolar spectroscopy: A practical guide. *J. Magn. Reson.* **2022**, *338*, 107186.
- (9) Jeschke, G.; Koch, A.; Jonas, U.; Godt, A. Direct Conversion of EPR Dipolar Time Evolution Data to Distance Distributions. *J. Magn. Reson.* **2002**, *155*, 75–82.
- (10) Chiang, Y.-W.; Borbat, P. P.; Freed, J. H. Maximum Entropy: A Complement to Tikhonov Regularization for Determination of Pair Distance Distributions by Pulsed ESR. *J. Magn. Reson.* **2005**, *177*, 184–196.
- (11) Jeschke, G.; Chechik, V.; Ionita, P.; Godt, A.; Zimmermann, H.; Banham, J.; Timmel, C. R.; Hilger, D.; Jung, H. DeerAnalysis2006—A Comprehensive Software Package for Analyzing Pulsed ELDOR Data. *Appl. Magn. Reson.* **2006**, *30*, 473–498.

951 (12) Sen, K. I.; Logan, T. M.; Fajer, P. G. Protein Dynamics and Monomer-Monomer
952 Interactions in AntR Activation by Electron Paramagnetic Resonance and Double Electron-Electron Resonance. *Biochem.*
953 **2009**, *46*, 11639–11649.

954 (13) Brandon, S.; Beth, A. H.; Hustedt, E. The Global Analysis of DEER Data. *J. Magn. Reson.* **2012**, *218*, 93–104.

955 (14) Stein, R. A.; Beth, A. H.; Hustedt, E. A Straightforward Approach to the Analysis of Double Electron–Electron Resonance Data. *Methods Enzymol.* **2015**, *563*, 531–567.

956 (15) Edwards, T. H.; Stoll, S. Optimal Tikhonov Regularization for DEER Spectroscopy. *J. Magn. Reson.* **2018**, *288*, 58–68.

957 (16) Sweger, S. R.; Pribitzer, S.; Stoll, S. Bayesian Probabilistic Analysis of DEER Spectroscopy Data using Parametric Distance Distribution Models. *J. Phys. Chem. A* **2020**, *124*, 6193–6202.

958 (17) Fábregas Ibáñez, L.; Jeschke, G.; Stoll, S. DeerLab: A Comprehensive Toolbox for Analyzing Dipolar EPR Spectroscopy Data. *Magn. Reson.* **2020**, *1*, 209–224.

959 (18) Hustedt, E. J.; Marinelli, F.; Stein, R. A.; Faraldo-Gómez, J. D.; Mchaourab, H. S. Confidence Analysis of DEER Data and its Structural Interpretation with Ensemble-Biased Metadynamics. *Biophys. J.* **2018**, *115*, 1200–1216.

960 (19) Bretthorst, G. L. Bayesian analysis. III. Applications to NMR signal detection, model selection, and parameter estimation. *J. Magn. Reson.* **1990**, *88*, 571–595.

961 (20) Hao, J.; Astle, W.; De Iorio, M.; Ebbels, T. M. D. BATMAN – an R package for the automated quantification of metabolites from nuclear magnetic resonance spectra using a Bayesian model. *Bioinformatics* **2012**, *28*, 2088–2090.

962 (21) Ikeya, T.; Ikeda, S.; Kigawa, T.; Ito, Y.; Güntert, P. Protein NMR structure refinement based on Bayesian inference. *J. Phys.: Conf. Ser.* **2016**, *699*, 012005.

963 (22) Edwards, T. H.; Stoll, S. A Bayesian Approach to Quantifying Uncertainty from Experimental Noise in DEER Spectroscopy. *J. Magn. Reson.* **2016**, *270*, 87–97.

964 (23) De Lorenzi, F.; Weinmann, T.; Bruderer, S.; Heitmann, B.; Henrici, A.; Stingelin, S. Bayesian analysis of 1D ¹H-NMR spectra. *J. Magn. Reson.* **2024**, *364*, 107723.

965 (24) Altenbach, C. LongDistances. <https://sites.google.com/site/altenbach/labview-programs/epr-programs/long-distances>.

966 (25) Fábregas-Ibáñez, L.; Tessmer, M. H.; Jeschke, G.; Stoll, S. Dipolar Pathways in Dipolar EPR Spectroscopy. *Phys. Chem. Chem. Phys.* **2022**, *24*, 2504–2520.

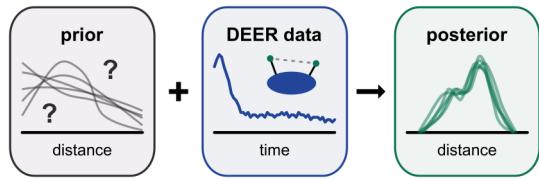
967 (26) Kattnig, D. R.; Reichenwallner, J.; Hinderberger, D. Modeling Excluded Volume Effects for the Faithful Description of the Background Signal in Double Electron–Electron Resonance. *J. Phys. Chem. B* **2013**, *117*, 16542–16557.

968 (27) Bardsley, J. M.; Fox, C. An MCMC Method for Uncertainty Quantification in Nonnegativity Constrained Inverse Problems. *Inverse Probl. Sci. Eng.* **2012**, *20*, 477–498.

969 (28) Bardsley, J. M.; Hansen, P. C. MCMC Algorithms for Computational UQ of Nonnegativity Constrained Linear Inverse Problems. *SIAM J. Sci. Comput.* **2020**, *42*, A1269–A1288.

970 (29) Robert, C. P.; Casella, G. *Monte Carlo Statistical Methods*; Springer: New York, 2010.

1037 (30) Bro, R.; De Jong, S. A fast non-negativity-1079
1038 constrained least squares algorithm. *J.1080
1039 Chemom.* **1997**, *11*, 393–401. 1081
1040 (31) Moré, J. J.; Toraldo, G. On the Solution of 1082
1041 Large Quadratic Programming Problems 1083
1042 with Bound Constraints. *Siam J. Control 1084
1043 Optim.* **1991**, *1*, 93–113. 1085
1044 (32) Van Benthem, M. H.; Keenan, M. R. Fast 1086
1045 algorithm for the solution of large-scale 1087
1046 non-negativity-constrained least squares 1088
1047 problems. *J. Chemom.* **2004**, *18*, 441–450. 1089
1048 (33) Hoffman, M. D.; Gelman, A. The No-U-1090
1049 Turn Sampler: Adaptively Setting Path 1091
1050 Lengths in Hamiltonian Monte Carlo. 1092
1051 *J. Mach. Learn. Res.* **2014**, *15*, 1593– 1093
1052 1623. 1094
1053 (34) Neal, R. M. *MCMC Using Hamiltonian 1096
1054 Dynamics*; CRC Press: Boca Raton, 2011. 1097
1055 (35) Betancourt, M. A Conceptual Introduc-1098
1056 tion to Hamiltonian Monte Carlo. 2018; 1099
1057 <https://arxiv.org/abs/1701.02434>. 1100
1058 (36) Salvatier, J.; Wiecki, T. V.; Fonnes-1102
1059 beck, C. Probabilistic Programming in1103
1060 Python Using PyMC3. *Peer J. Com-1104
1061 put. Sci.* **2016**, *2*, e55. 1105
1062 (37) Gelman, A.; Rubin, D. B. Inference from1106
1063 Iterative Simulation Using Multiple Se-1107
1064 quences. *Stat. Sci.* **1992**, *7*, 457–472. 1108
1065 (38) Brooks, S. P.; Gelman, A. General Meth-1109
1066 ods for Monitoring Convergence of Itera-1110
1067 tive Simulations. *J. Comput. Graph. Stat.*
1068 **1998**, *7*, 434–455. 1105
1069 (39) Vehtari, A.; Gelman, A.; Simpson, D.;
1070 Carpenter, B.; Bürkner, P.-C. Rank-
1071 Normalization, Folding, and Localization:
1072 An Improved \hat{R} for Assessing Convergence
1073 of MCMC. *Bayesian Anal.* **2021**, *1*, 667–
1074 718. 1109
1075 (40) Engelhardt, B. E.; Adams, R. P.
1076 Bayesian Structured Sparsity from
1077 Gaussian Fields. 2014; <https://arxiv.org/abs/1407.2235>. 1105
1078 (41) Edwards, T. H.; Stoll, S. Synthetic
1079 Test Data Set for DEER Spectroscopy
1080 Based on T4 Lysozyme. DOI:
1081 10.6069/H5S75DCG, 2018;
1082 10.6069/H5S75DCG.
1083 (42) Fábregas-Ibáñez, L.; Jeschke, G.; Stoll, S.
1084 Compactness regularization in the
1085 analysis of dipolar EPR spectroscopy.
1086 *J. Magn. Reson.* **2022**, *339*, 107218.
1087 (43) Schiemann, O.; Heubach, C. A.; Abdullin, D.; Ackermann, K.; Azarkh, M.; Bagryanskaya, E. G.; Drescher, M.; Endeward, B.; Freed, J. H.; Galazzo, L. et al. Benchmark test and guidelines for DEER/PELDOR experiments on nitroxide-labeled biomolecules. *J. Am. Chem. Soc.* **2021**, *143*, 17875–17890.
1088 (44) Peter, M. F.; Tuukkanen, A. T.; Heubach, C. A.; Selsam, A.; Duthie, F. G.; Svergun, D. I.; Schiemann, O.; Hagelueken, G. Studying conformational changes of the Yersinia type-III-secretion effector YopO in solution by integrative structural biology. *Structure* **2019**, *27*, 1416–1426.
1089 (45) Dzuba, S. A. The determination of pair-distance distribution by electron–electron resonance: regularization by the length of distance discretization with Monte Carlo calculations. *J. Magn. Reson.* **2016**, *269*, 113–119.



1112

1113