PHYSICAL REVIEW E **98**, 043308 (2018)

Numerical inversion methods for recovering negative amplitudes in two-dimensional nuclear magnetic resonance relaxation-time correlations

Thusara C. Chandrasekera
Churchill College, University of Cambridge, Cambridge CB3 0DS, United Kingdom

Jonathan Mitchell[*]
Schlumberger Cambridge Research, High Cross, Madingley Road, Cambridge CB3 0EL, United Kingdom

(Received 30 July 2018; published 25 October 2018)

Two-dimensional nuclear magnetic resonance measurements are ubiquitous in the literature, with correlations of longitudinal T_1 and transverse T_2 relaxation times used extensively to characterize porous media. Decomposition of the signal acquired in the time domain to a pseudocontinuous distribution of relaxation times is achieved using numerical inversion. A popular technique to generate a stable solution to this ill-posed problem in the presence of noise is Tikhonov regularization with a non-negativity constraint imposed on the output. However, coupling of the longitudinal and transverse eigenfunctions can generate eigenvalue pairs with apparent $T_2 > T_1$ and negative amplitude. Such apparent signal components are encountered in the classic example of Brownstein-Tarr "slow" diffusion in an isolated pore, and in weakly coupled pores governed by different relaxation rates. We show that when negative-amplitude components comprise $\geqslant 1\%$ of the total signal, the solution achieved by non-negative Tikhonov regularization is sufficiently distorted to prevent robust interpretation. We demonstrate two alternative inversion methods that recover the negative-amplitude components: (1) half-bound Tikhonov regularization assigns a negative amplitude to any peak with apparent $T_2 > T_1$, and (2) the optimization problem is expressed as a ℓ_2 regression with ℓ_1 penalization and a solution estimated using a primal-dual algorithm without constraint on the output sign. These methods are applicable to T_1-T_2 experiments on porous materials characterized by a hierarchy of length scales, such as biological cells, cement, and limestone.

DOI: 10.1103/PhysRevE.98.043308

I. INTRODUCTION

Nuclear magnetic resonance (NMR) relaxation-time distributions offer a powerful tool for materials characterization. Longitudinal T_1 and transverse T_2 relaxation rates are used in lieu of chemical sensitivity at low magnetic field strengths [1]. The two-dimensional (2D) T_1-T_2 correlation was introduced for the study of petroleum reservoir rocks [2] to provide improved fluid-phase contrast and an estimator of wettability [3], and detailed interpretation of unconventional resources containing solidlike hydrocarbons [4,5]. The experiment has been applied to a diverse range of materials including cements [6], dairy products [7], and heterogeneous catalysts [8]. In bulk liquids, it is generally true that $T_1 = T_2$. However, these relaxation times are modified for liquids imbibed in porous media due to surface adsorption (reduced mobility) and interactions with surface relaxation sinks [9,10]. Under these conditions, $T_1 \geqslant T_2$ with the ratio T_1/T_2 considered proportional to the strength of surface adsorption [3,11].

Two-dimensional relaxation-time correlations have been used to study chemical or diffusive exchange [12,13]. Some of these experiments, such as T_2-T_2 [14] or T_1-T_1 [15], employ an explicit storage interval to encode the exchange rate. Off-diagonal relaxation-time components are observed in the correlation plots whenever $T_1 \neq T_2$. The amplitude of these "exchange peaks" as a function of storage time is fitted to estimate the coupling rate constant for spins in different environments [16,17]. The T_2-T_2 correlation is preferred because the amplitude of the distribution is entirely positive under practical conditions [17]. Similar off-diagonal components are present in T_1-T_2 correlations, despite the absence of a defined storage interval, and components of the distribution with apparent $T_2 > T_1$ have negative amplitude [18]. The off-diagonal components evolve as a mathematical consequence of projecting the eigenfunction basis for longitudinal relaxation on to the eigenfunction basis for transverse relaxation, or vice versa [18,19].

The relaxation rates for diffusing spins are governed by the Bloch-Torrey equation [20]. This partial differential equation can be expressed in terms of a scaled magnetization m in separate equations for longitudinal and transverse relaxation. The version for longitudinal relaxation is

$$\frac{\partial}{\partial t}[m_1(\mathbf{r},t) - m^{\mathrm{eq}}] = \mathcal{D}\nabla^2[m_1(\mathbf{r},t) - m^{\mathrm{eq}}] - \frac{m_1(\mathbf{r},t) - m^{\mathrm{eq}}}{T_1^{\mathrm{b}}}, \qquad (1)$$

where $m_1(\mathbf{r},t)$ is scaled in proportion to the longitudinal component of the magnetization M_z, m^{eq} is similarly scaled in proportion to the equilibrium magnetization M^{eq}, and $1/T_1^{\mathrm{b}}$ is the bulk longitudinal relaxation rate. The required Robin boundary condition for the magnetization of a liquid, imbibed

---

[*] Corresponding author: jmitchell16@slb.com

in a porous material, is

$$\hat{\mathbf{n}}\cdot\mathcal{D}\nabla[m_1(\mathbf{r},t)-m^{\mathrm{eq}}]+\rho_1[m_1(\mathbf{r},t)-m^{\mathrm{eq}}]=0, \tag{2}$$

where $\hat{\mathbf{n}}$ is the unit outward normal vector for the boundary and ρ_1 is the longitudinal surface relaxivity. For transverse relaxation, the Bloch-Torrey equation is written in the familiar form

$$\frac{\partial}{\partial t}m_2(\mathbf{r},t)=\mathcal{D}\nabla^2 m_2(\mathbf{r},t)-\frac{m_2(\mathbf{r},t)}{T_2^{\mathrm{b}}}, \tag{3}$$

where $m_2(\mathbf{r},t)$ is scaled in proportion to the magnitude of the transverse magnetization (or either component on resonance in the rotating frame, e.g., M_x) and $1/T_2^{\mathrm{b}}$ is the bulk transverse relaxation rate. The Robin boundary condition in this case is

$$\hat{\mathbf{n}}\cdot\mathcal{D}\nabla m_2(\mathbf{r},t)+\rho_2 m_2(\mathbf{r},t)=0, \tag{4}$$

where ρ_2 is the transverse surface relaxivity.

The T_1-T_2 correlation experiment comprises of a period of longitudinal magnetization preparation by inversion recovery [21] with delay t_1 followed by a period of magnetization detection during a Carr-Purcell-Meiboom-Gill (CPMG) sequence [22,23] in which n is the number of spin echoes separated by echo time t_{e}. The equilibrium magnetization m^{eq} is scaled such that $\int m^{\mathrm{eq}}\,dV(\mathbf{r})=1$. The integral is over the volume of the entire geometry; m^{eq} will of course be the reciprocal of the volume that can be detected by NMR, given this normalization to unit volume integral. Perfect inversion of the equilibrium magnetization is assumed for the inversion recovery part of the pulse sequence, so the initial longitudinal value $m_1(\mathbf{r},0)=-m^{\mathrm{eq}}$ is scaled such that $\int m_1(\mathbf{r},0)\,dV(\mathbf{r})=-1$. The required solution to Eqs. (1) and (2) for the longitudinal magnetization is obtained by eigenfunction expansion. Each longitudinal relaxation-diffusion eigenfunction $\psi_j(\mathbf{r})$ satisfies $(-\mathcal{D}\nabla^2+1/T_1^{\mathrm{b}})\psi_j(\mathbf{r})=\psi_j(\mathbf{r})/T_{1,j}^{\mathrm{app}}$ and $\hat{\mathbf{n}}\cdot\mathcal{D}\nabla\psi_j(\mathbf{r})+\rho_1\psi_j(\mathbf{r})=0$ such that the eigenvalue (apparent longitudinal relaxation rate) is $1/T_{1,j}^{\mathrm{app}}$. These eigenfunctions are defined to be orthonormal, viz., $\int\psi_i^*(\mathbf{r})\psi_j(\mathbf{r})\,dV(\mathbf{r})=\delta_{ij}$ (where the Krönecker delta has the value of 1 for $i=j$ and 0 for $i\neq j$). The eigenvalue-eigenfunction pairs are indexed in order of increasing magnitude of eigenvalue. The scaled magnetization due to longitudinal relaxation during time t_1 following perfect inversion of the equilibrium magnetization is

$$m_1(\mathbf{r},t_1)=\sum_{j=0}^{\infty}\psi_j(\mathbf{r})\left[1-2\exp\left\{-\frac{t_1}{T_{1,j}^{\mathrm{app}}}\right\}\right]a_j, \tag{5}$$

where $a_j=\int\psi_j^*(\mathbf{r})m^{\mathrm{eq}}\,dV(\mathbf{r})$ are the equilibrium amplitudes for the eigenfunctions $\psi_j(\mathbf{r})$. The expression for $m_1(\mathbf{r},t_1)$ in Eq. (5) is regarded as a generalized Fourier series with the eigenfunctions $\psi_j(\mathbf{r})$ taking the place of the usual sinusoidal Fourier basis and the eigenvalues generalize the angular frequency. The discrete eigenvalues describe the spectrum of the longitudinal relaxation-diffusion operator $(-\mathcal{D}\nabla^2+1/T_1^{\mathrm{b}})$ with the given boundary condition. In ascending order of eigenvalue magnitude, the eigenfunctions exhibit more zero crossings where they transition between positive and negative amplitude, similar to the Fourier basis.

The longitudinal magnetization is then rotated to the transverse plane for a time $t_2=nt_{\mathrm{e}}$. The required solution to Eqs. (3) and (4) for the transverse magnetization is obtained by another eigenfunction expansion. Each transverse relaxation-diffusion eigenfunction $\phi_i(\mathbf{r})$ satisfies $(-\mathcal{D}\nabla^2+1/T_{2,\mathrm{b}})\phi_i(\mathbf{r})=\phi_i(\mathbf{r})/T_{2,i}^{\mathrm{app}}$ and $\hat{\mathbf{n}}\cdot\mathcal{D}\nabla\phi_i(\mathbf{r})+\rho_2\phi_i(\mathbf{r})=0$ such that the eigenvalue (apparent transverse relaxation rate) is $1/T_{2,i}^{\mathrm{app}}$. These eigenfunctions are also defined to be orthonormal, viz., $\int\phi_i^*(\mathbf{r})\phi_j(\mathbf{r})\,dV(\mathbf{r})=\delta_{ij}$. The eigenvalue-eigenfunction pairs are indexed in order of increasing magnitude of eigenvalue. The scaled magnetization due to transverse relaxation during the time t_2 is thus

$$m_2(\mathbf{r},t_1,t_2)=\sum_{i=0}^{\infty}\phi_i(\mathbf{r})\exp\left\{-\frac{t_2}{T_{2,i}^{\mathrm{app}}}\right\}b_i, \tag{6}$$

where $b_i=\int\phi_i^*(\mathbf{r})m_1(\mathbf{r},t_1)\,dV(\mathbf{r})$. The observed signal $h(t_1,t_2)\propto\int m_2(\mathbf{r},t_1,t_2)\,dV(\mathbf{r})$.

The eigenvalues and eigenfunctions are determined by the pore geometry and surface relaxivity for each case of longitudinal and transverse relaxation. The eigenfunctions are the same as the diffusion "eigenmodes" (Laplacian eigenfunctions) that feature in [18], but the corresponding eigenvalues are increased from the diffusion case by the bulk relaxation rate (e.g., $1/T_1^{\mathrm{b}}$). The required eigenfunctions and eigenvalues can be determined analytically for simple geometries [24]. A method for calculating the expected signal contributions in the T_1-T_2 correlation, from the possible combinations of the "longitudinal" and "transverse" eigenvalues, is given for such simple geometries in [18]. The result can be expressed in terms of the amplitudes of the peaks positioned at apparent relaxation times $\{T_{1,j}^{\mathrm{app}},T_{2,i}^{\mathrm{app}}\}$ by

$$f(T_1,T_2)=\sum_{i,j}d_i\,\delta\left(\log_{10}\frac{T_2}{T_{2,i}^{\mathrm{app}}}\right)Q_{ij}\,\delta\left(\log_{10}\frac{T_1}{T_{1,j}^{\mathrm{app}}}\right)a_j, \tag{7}$$

where $d_i=\int\phi_i(\mathbf{r})\,dV(\mathbf{r})$, $Q_{ij}=\int\phi_i^*(\mathbf{r})\psi_j(\mathbf{r})\,dV(\mathbf{r})$, and a_j is as above in Eq. (5). The density distribution $f(T_1,T_2)$ is normalized such that $\int_{T_2=0}^{\infty}\int_{T_1=0}^{\infty}f(T_1,T_2)\,d(\log_{10}T_1)\,d(\log_{10}T_2)=1$. The presence of the Dirac delta factors in Eq. (7) mean that $f(T_1,T_2)$ is sparse. The peak integral over $\log_{10}T_1$ and $\log_{10}T_2$ for each $\{T_{1,j}^{\mathrm{app}},T_{2,i}^{\mathrm{app}}\}$ contribution is $A_{ij}=d_iQ_{ij}a_j$. Each peak thus corresponds to a different combination of eigenvalues in location and the amplitude of the peak depends on the corresponding eigenfunctions.

In multidimensional relaxation-time experiments such as the T_1-T_2 correlation analyzed above, the eigenfunctions relevant for one measurement interval (ψ_j, $j=0,1,2,\ldots$, having corresponding amplitudes a_j and eigenvalues $1/T_{1,j}^{\mathrm{app}}$) are projected on to the eigenfunction basis ϕ_i relevant for the next measurement interval. That is, the generalized Fourier series based on $\psi_j(\mathbf{r})$ in Eq. (5) is reexpressed as another generalized Fourier series based on $\phi_i(\mathbf{r})$ such that $m_2(\mathbf{r},0)=m_1(\mathbf{r},t_1)$. As the two eigenfunction bases have different spatial dependence over the same volume, this projection is nontrivial and all combinations of $\{\psi_j,\phi_i\}$ are nonorthogonal (unless $\rho_1=\rho_2$) yielding nonzero Q_{ij} for $i\neq j$. Consequently, we expect to find a nonzero signal contribution from each eigenvalue combination $\{1/T_{1,j}^{\mathrm{app}},1/T_{2,i}^{\mathrm{app}}\}$, although most "high-index" combinations (e.g., i, $j>2$) have negligible amplitude relative to the “low-index” combinations. The combinations of eigenvalues with $i = j$ give the “diagonal” peaks in a plot of $f(T_1, T_2)$ (i.e., near the $T_1 = T_2$ diagonal) and the combinations of eigenvalues with $i \neq j$ give the “off-diagonal” peaks. Note that the plot is always asymmetric in peak locations about $T_1 = T_2$ except in the limiting case of $T_{1,j}^{\text{app}} = T_{2,i}^{\text{app}}$. Some eigenvalue combinations, typically with $T_{1,j}^{\text{app}} < T_{2,i}^{\text{app}}$ if $T_1^{\text{b}} \approx T_2^{\text{b}}$, give peaks that exhibit negative amplitude [18]. The presence of negative amplitudes mean that $f(T_1, T_2)$ cannot be regarded as a probability density function in this situation.

The observed signal $h(t_1, t_2) \propto \int m_2(\mathbf{r}, t_1, t_2)\, dV(\mathbf{r})$, and the density distribution of the time constants $f(T_1, T_2)$, are related by a Fredholm integral equation of the first kind:

$$\frac{h(t_1, t_2)}{h(0, 0)} = \int_{T_2=0}^{\infty} \int_{T_1=0}^{\infty} k_1(T_1, t_1)\, k_2(T_2, t_2)\, f(T_1, T_2) \times d(\log_{10} T_1)\, d(\log_{10} T_2), \tag{8}$$

where the subscripts $\{1, 2\}$ serve to denote the experiment dimensions, with the integration kernels $k_1(T_1, t_1) = 1 - 2\exp\{-t_1/T_1\}$ and $k_2(T_2, t_2) = \exp\{-t_2/T_2\}$. The overall integration kernel is separable in this case, viz., $k_0(T_1, t_1, T_2, t_2) = k_1(T_1, t_1)\, k_2(T_2, t_2)$. Estimates of $f(T_1, T_2)$ are obtained by solving this inverse problem. The expression $h(t_1, t_2)/h(0, 0)$ is the normalized signal, which has a maximum value of unity [given k_0 and the normalization of $f(T_1, T_2)$ imposed here]. To test reconstruction of $f(T_1, T_2)$ via simulation in geometries with known eigenvalues and eigenfunctions, an expression for the expected normalized signal could be obtained from the right-hand side of Eq. (7) simply by replacing the pair of Dirac delta factors with $k_0(T_{1,j}^{\text{app}}, t_1, T_{2,i}^{\text{app}}, t_2)$, for example. However, measured data $g(t_1, t_2)$ are comprised of signal $h(t_1, t_2)$ and noise $e(t_1, t_2)$ such that $g(t_1, t_2) = h(t_1, t_2) + e(t_1, t_2)$. The overall integration kernel k_0 is based on exponential functions, which typically means that the reconstructed estimate $\hat{f}(T_1, T_2)$ is extremely sensitive to the noise $e(t_1, t_2)$, and may contain a nonsensical distribution. The inverse problem is thus ill posed and regularization is required to achieve a solution $\hat{f}(T_1, T_2)$ that is stable to small changes in the measured data. A popular method for solving Eq. (8) is to formulate the optimization problem as a ℓ_2-penalized regression. Tikhonov regularization [25] provides an optimum smooth solution when constrained to be non-negative. A computationally efficient method for solving this optimization problem in 2D when the kernel is separable was introduced in [26], based on [27], and generalized to nonseparable kernels in [28]. To achieve a stable and sensible solution, it is necessary to enforce the non-negativity constraint on $\hat{f}(T_1, T_2)$ as part of the regularization. This condition precludes the detection of negative-amplitude components in the resulting distribution [18,29]. Negative-amplitude components have been observed in time-domain data [30,31], but cannot be recovered in 2D distributions using the method of [26].

Recently, there has been interest in applying an alternative numerical inversion method to NMR data. The least absolute shrinkage and selection operator (LASSO) [32,33] combines ℓ_2 regression with ℓ_1 penalization to balance fidelity to the signal measurements against sensitivity of the density distribution estimate to noise in the measurements. Several algorithms are available for solving the LASSO. Iterative soft thresholding is a popular method for addressing this class of linear inverse problem (see [34] and references therein). For example, Stern *et al.* obtained chemical spectra from truncated free induction decays (FIDs) without artifacts [35]. This algorithm tends towards a sparse solution (rather than a smooth, continuous solution) by setting low-amplitude components to zero. Consequently, it is well suited to recovering chemical spectra expected to contain signal components at discrete frequencies. The fast iterative shrinkage-thresholding algorithm (FISTA) of Beck and Teboulle [36] improves the computational efficiency of the numerical inversion. This primal algorithm, which alternates steepest descent for ℓ_2 regression with soft thresholding for ℓ_1 penalization, has been applied to 2D NMR relaxation-time correlations [37]. An alternative approach is the primal-dual algorithm of Chambolle and Pock [38], adapted by Reci *et al.* [39,40] to balance sparsity and smoothness in the solution. The ℓ_1 penalty parts of the primal objective function are expressed as convex conjugates of their dual versions, introducing extra variables to estimate in the fit. The resulting saddle-point problem is solved by a proximal splitting method [38]. Alternating between ℓ_2 smoothing on the primal variables and ℓ_∞ pruning on the dual variables implicitly implements net ℓ_1-norm penalties on the primal variables; i.e., the primal objective function is indirectly minimized with respect to its primal variables. Recent implementations retained a non-negativity constraint for the solution [39,40], although there is no necessity to do so.

In Sec. II we summarize the Tikhonov regularization and primal-dual algorithms for solving the inverse problem. Generation of simulated time-domain NMR data containing relaxation contributions of negative amplitude is described in Sec. III. The inversion techniques are applied to these simulated data in Sec. IV and a practical application to experimental data is presented in Sec. V. We discuss the advantages and limitations of each method and offer general guidance for selecting an appropriate numerical inversion technique for application to experimental data.

II. NUMERICAL INVERSIONS

To describe the inversion methods, we adapt our previous notation [28] where necessary. Variables, continuous functions, and scalars are denoted by *italics*. Vector arrays are denoted by lower case bold. Matrix arrays are denoted by upper case bold. The ℓ_1 norm of a vector array is indicated by $\|\ldots\|_1$, which is the sum of absolute element values. The ℓ_2 norm of a vector array is indicated by $\|\ldots\|_2$ and the related Frobenius norm of a matrix array is indicated by $\|\ldots\|_F$; both are the root sum of squared element values. The ℓ_∞ norm of a vector array is indicated by $\|\ldots\|_\infty$, which is the maximum of absolute element values. The Krönecker product (nested multiplication) for matrices is denoted by $\otimes$ and the Hadamard product (element-wise multiplication) for vectors or matrices is denoted by $\odot$. All calculations were implemented in the MATLAB(Mathworks, Inc.) programming environment on a desktop PC.

To provide simple demonstrations of the inversion methods, we consider a one-dimensional (1D) Fredholm integral equation

$$\frac{h(t)}{h(0)} = \int_{T=0}^{\infty} k_0(T,t) f(T) d(\log_{10} T), \tag{9}$$

and assume an exponential kernel function $k_0(T,t) = \exp(-t/T)$. This case essentially describes the 1D T_2 experiment, i.e., $k_0(T,t) \equiv k_2(T,t)$. Note that T may be encoded as a recovery or decay curve in experimental design [15,21,41] without consequence for the analysis. Also, there is no requirement for the kernel function to be exponential; examples of nonexponential kernels are found in [42,43]. Practically, 1D distributions of T_2 derived from experimental CPMG data will be positive always (assuming equilibrium magnetization at $t = 0$). However, for the example studied in this section, $f(T)$ is allowed to take negative values.

We wish to determine an optimum solution $\hat{f}$ that recovers the signal h from noisy measurements g. The number of measurements made is q, each corresponding to a different relaxation delay in the NMR experiment. As reviewed in [28], integrals involving a given kernel function $k_0(T,t)$ and distribution f are simply approximated by a discrete sum for numerical calculations. The distribution is represented by a $p \times 1$ column vector $\mathbf{x}$, each element an amplitude corresponding to a different prospective time constant T. Typically, $p < q$ is chosen for inversions for NMR relaxation experiments, albeit the requirement is not essential with sufficient regularization in the inversion method. Normalization of the distribution is ignored during calculations since the discrete approximation of f can always be obtained from $\mathbf{x}$ by rescaling the latter to yield unit integral over $\log_{10} T$. The kernel function is discretely sampled forming the elements of a $q \times p$ matrix $\mathbf{K}_0$ where each row corresponds to a different measurement delay and each column corresponds to an element of $\mathbf{x}$. The approximated integral is thus given by the vector-matrix product $\mathbf{K}_0\mathbf{x}$. The measured NMR data are represented in a $q \times 1$ column vector $\mathbf{g}$ such that

$$\mathbf{g} = \mathbf{K}_0\mathbf{x} + \mathbf{e}, \tag{10}$$

where $\mathbf{e}$ is a $q \times 1$ column vector in which each element is the noise contribution to the corresponding element in $\mathbf{g}$. An example signal $\mathbf{g}$ is illustrated in Fig. 1(a) (black solid line). It consists of $q = 128$ points spaced evenly across the logarithmic scale shown. The noise $e(t)$ is simulated with a normal distribution of mean zero and variance set to give a signal-to-noise ratio SNR ≈ 1000. The simulated signal is entirely positive, but described by three discrete rates $1/T = 45$, 10, and 2 s^{-1} with amplitudes -2, 1, and 2, respectively. The net signal is thus $g(t) = -2\exp\{-45\ \text{s}^{-1}\ t\} + \exp\{-10\ \text{s}^{-1}\ t\} + 2\exp\{-2\ \text{s}^{-1}\ t\} + e(t)$ and the true density distribution is $f(T) = -2\,\delta(\log_{10}\{45\ \text{s}^{-1}\ T\}) + \delta(\log_{10}\{10\ \text{s}^{-1}\ T\}) + 2\,\delta(\log_{10}\{2\ \text{s}^{-1}\ T\})$. The three Dirac delta terms are not appropriate for plotting due to their infinite extent. A broadened version of $f(T)$ is instead shown in Fig. 1(b). The distribution plotted corresponds to the discrete elements of vector $\mathbf{x}$ with the $p = 64$ points spaced evenly across the logarithmic scale shown. This vector would be the ideal result of any reconstruction method using the same discrete sampling for the columns of matrix $\mathbf{K}_0$ in Eq. (10). The true distribution is most accurately depicted as a cumulative integral $F(T) = \int_{T'=0}^{T} f(T') d(\log_{10} T')$ in Fig. 1(c).

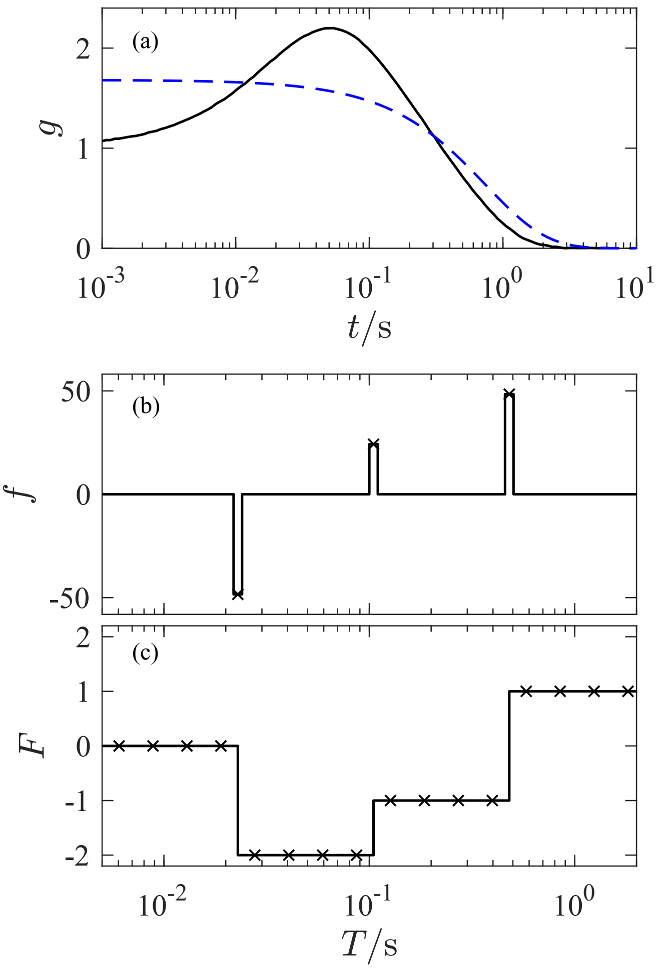

FIG. 1. Simulated data used to demonstrate the inversion methods presented here. The signal $\mathbf{g}$ is shown in (a), black solid line. The true relaxation-time distribution $f(T)$ contains three discrete relaxation rates with positive or negative amplitude: an approximate version is shown in (b). The true distribution is given in (c) as a cumulative integral $F(T) = \int_{T'=0}^{T} f(T') d(\log_{10} T')$. The black × symbols indicate points reproduced in subsequent figures for comparison. The fitted data $\hat{\mathbf{g}} = \mathbf{K}_0\hat{\mathbf{x}}$ obtained by non-negative Tikhonov regularization (Sec. II A) are given in (a), blue dashed line. The fitted data obtained by unconstrained Tikhonov regularization, half-bound Tikhonov regularization (Sec. II B), and the primal-dual algorithm (Sec. II C) are not shown due to being indistinguishable from the original signal $\mathbf{g}$ at the scale plotted. See online version for color.

The singular value decomposition (SVD) of a matrix is a tool used commonly in numerical calculations for inverse problems [44, Chap. 3]. The relevant version for the inversions here is $\mathbf{K}_0 = \mathbf{U}_0\mathbf{S}_0\mathbf{V}_0^{\text{T}}$. The $q \times p$ matrix $\mathbf{U}_0$ has orthonormal columns so $\mathbf{U}_0^{\text{T}}\mathbf{U}_0$ yields a $p \times p$ identity matrix. The $p \times p$ matrix $\mathbf{S}_0$ is diagonal and contains the singular values of $\mathbf{K}_0$, $s_{0,i} = S_{0,ii}$, with largest value $s_{0,\text{max}}$ and smallest nonzero value $s_{0,\text{min}}$. The $p \times p$ matrix $\mathbf{V}_0$ is orthogonal so $\mathbf{V}_0^{\text{T}}\mathbf{V}_0$ yields a $p \times p$ identity matrix.

The measurement equation (10) is also applicable to 2D data [26], viz., Eq. (8) for the T_1-T_2 correlation experiment.

The distribution $f(T_1, T_2)$ is represented in non-normalized form by a $p_2 \times p_1$ matrix $\mathbf{X}$, each row of which corresponds to a prospective value of T_2 and each column corresponds to a prospective value of T_1. The column vector $\mathbf{x}$ is formed by stacking the columns of $\mathbf{X}$ below each other, viz., $\mathbf{x} = \mathrm{vec}(\mathbf{X})$, so $p = p_1 p_2$. The measured 2D NMR data are represented by a $q_2 \times q_1$ matrix $\mathbf{G}$, each row of which corresponds to a CPMG delay t_2 (the direct dimension) and each column corresponds to a longitudinal relaxation delay t_1 (the indirect dimension). The column vector $\mathbf{g}$ is formed by stacking the columns of $\mathbf{G}$, viz., $\mathbf{g} = \mathrm{vec}(\mathbf{G})$, so $q = q_1 q_2$. The kernel function for transverse relaxation $k_2(T_2, t_2)$ is sampled to form a $q_2 \times p_2$ matrix $\mathbf{K}_2$, similar to the matrix $\mathbf{K}_0$ for the above 1D example. A separate $q_1 \times p_1$ matrix $\mathbf{K}_1$ is formed for the longitudinal relaxation kernel function $k_1(T_1, t_1)$. The $q \times p$ matrix $\mathbf{K}_0$ for this separable 2D case is generated by the Krönecker product $\mathbf{K}_0 = \mathbf{K}_1 \otimes \mathbf{K}_2$. This matrix corresponds to the discrete sampling of the overall kernel function $k_0(T_1, t_1, T_2, t_2) = k_1(T_1, t_1)\, k_2(T_2, t_2)$. The individual SVDs $\mathbf{K}_1 = \mathbf{U}_1\mathbf{S}_1\mathbf{V}_1^{\mathrm{T}}$ and $\mathbf{K}_2 = \mathbf{U}_2\mathbf{S}_2\mathbf{V}_2^{\mathrm{T}}$ can be used to generate $\mathbf{U}_0 = \mathbf{U}_1 \otimes \mathbf{U}_2$, $\mathbf{S}_0 = \mathbf{S}_1 \otimes \mathbf{S}_2$, and $\mathbf{V}_0 = \mathbf{V}_1 \otimes \mathbf{V}_2$ [26]. For clarity, the factored matrix sizes here are $q_1 \times p_1$ for $\mathbf{U}_1$, $p_1 \times p_1$ for $\mathbf{S}_1$, $p_1 \times p_1$ for $\mathbf{V}_1$, $q_2 \times p_2$ for $\mathbf{U}_2$, $p_2 \times p_2$ for $\mathbf{S}_2$, and $p_2 \times p_2$ for $\mathbf{V}_2$. Furthermore, the singular values $s_{1,j} = S_{1,jj}$, $s_{2,i} = S_{2,ii}$ have respective maximum values $s_{1,\max}$, $s_{2,\max}$ and respective minimum values $s_{1,\min}$, $s_{2,\min}$, so $s_{0,\max} = s_{1,\max}s_{2,\max}$ and $s_{0,\min} = s_{1,\min}s_{2,\min}$ for the T_1-T_2 correlation.

A. Non-negative Tikhonov regularization

A traditional method for solving linear inverse problems, the Tikhonov regularization [25], reaches an optimum smooth solution to the unconstrained ℓ_2-penalized regression (following [44, Chap. 8])

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \left\{ \|\mathbf{g} - \mathbf{K}_0\mathbf{x}\|_2^2 + \lambda^2 \|\mathbf{L}_0\mathbf{x}\|_2^2 \right\}, \tag{11}$$

where λ^2 is a tuning parameter, and $\mathbf{L}_0$ is either the identity matrix or a second finite-difference matrix (finite-difference Laplacian matrix) with p columns that dictates the nature of smoothing. A simple second finite-difference matrix can be generated, for example, as a $p \times p$ Toeplitz matrix with -2 on the main diagonal and 1 on each of the adjacent subdiagonals and superdiagonals (the matrix bandwidth is 3). If $\mathbf{L}_0$ is instead taken as the $p \times p$ identity matrix, the problem can be solved conveniently by SVD [44, Chap. 4]:

$$\begin{aligned}\hat{\mathbf{x}} &= \arg\min_{\mathbf{x}} \left\{ \|\mathbf{g} - \mathbf{K}_0\mathbf{x}\|_2^2 + \lambda^2 \|\mathbf{x}\|_2^2 \right\} \\ &= \mathbf{V}_0\big(\mathbf{w}_0 \odot \big(\mathbf{U}_0^{\mathrm{T}}\mathbf{g}\big)\big),\end{aligned} \tag{12}$$

where the elements of $p \times 1$ vector $\mathbf{w}_0$ are $w_{0,i} = s_{0,i}/(s_{0,i}^2 + \lambda^2)$. This method yields a solution for $\hat{\mathbf{x}}$ ($p \times 1$) that is biased to be shrunk towards zero and is stable to small changes in the measured data $\mathbf{g}$ ($q \times 1$) for sufficiently high values of the tuning parameter λ^2. When the value of λ^2 is too high, the solution is shrunk so far towards zero that it cannot represent the measured data faithfully. Methods for optimizing the value of λ^2 are discussed in [44, Chap. 5]. Even with a suitable choice of λ^2, however, this traditional form of Tikhonov regularization typically does not yield realistic results for NMR relaxometry problems. The smooth solutions are often too broad and undulate from positive to negative amplitude across the range of T, even in common situations where non-negative amplitudes are expected. A more sophisticated version of the technique is necessary, especially for NMR relaxometry experiments in more than one dimension [26].

Non-negative Tikhonov regularization aims to find more realistic solutions via the lower-bound-constrained ℓ_2-penalized regression

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}\geqslant \mathbf{0}} \left\{ \|\mathbf{g} - \mathbf{K}_0\mathbf{x}\|_2^2 + \lambda^2 \|\mathbf{L}_0\mathbf{x}\|_2^2 \right\}. \tag{13}$$

Note this method is often referred to incorrectly as an inverse Laplace transform (ILT) in the literature [45]. Regularization is imposed from three contributions of prior knowledge on f (and hence $\mathbf{x}$):

(1) the range (bandwidth) of $\log_{10} T$ is restricted,
(2) the distribution is biased to be smooth,
(3) the distribution is non-negative, $f(T) \geqslant 0$.

The size of the optimization problem can be reduced by projecting the data on a truncated singular value basis without significant change to the result [26]. Only signal components corresponding to singular values above a threshold determined by the SNR of the data are retained, leaving $\tilde{p}$ significant singular values. Quantities projected on to this truncated basis are indicated by a ˜ (tilde). The columns of $\mathbf{U}_0$ and $\mathbf{V}_0$ that are retained in $\tilde{\mathbf{U}}_0$ ($q \times \tilde{p}$) and $\tilde{\mathbf{V}}_0$ ($p \times \tilde{p}$), respectively, are those that correspond to the $\tilde{p}$ singular values retained in the diagonal of the reduced matrix $\tilde{\mathbf{S}}_0$ ($\tilde{p} \times \tilde{p}$). Setting $\tilde{\mathbf{g}} = \tilde{\mathbf{U}}_0^{\mathrm{T}}\mathbf{g}$ ($\tilde{p} \times 1$) and $\tilde{\mathbf{K}}_0 = \tilde{\mathbf{S}}_0\tilde{\mathbf{V}}_0^{\mathrm{T}}$ ($\tilde{p} \times p$), as well as taking $\mathbf{L}_0$ as the $p \times p$ identity matrix, the minimization in Eq. (13) is approximately reformulated as the following equivalent optimization problems:

$$\begin{aligned}&\min_{\mathbf{x}\geqslant 0} \left\{ \big\|\tilde{\mathbf{g}} - \tilde{\mathbf{K}}_0\mathbf{x}\big\|_2^2 + \lambda^2 \|\mathbf{x}\|_2^2 \right\} \\ &\quad= \lambda^2 \min_{\mathbf{x}\geqslant 0} \max_{\mathbf{c}} \left\{ -2\mathbf{x}^{\mathrm{T}}\tilde{\mathbf{K}}_0^{\mathrm{T}}\mathbf{c} - \big(\lambda^2\|\mathbf{c}\|_2^2 - 2\tilde{\mathbf{g}}^{\mathrm{T}}\mathbf{c}\big) + \|\mathbf{x}\|_2^2 \right\} \\ &\quad= \lambda^2 \max_{\mathbf{c}} \left\{ -\|\hat{\mathbf{x}}(\mathbf{c})\|_2^2 - \lambda^2\|\mathbf{c}\|_2^2 + 2\tilde{\mathbf{g}}^{\mathrm{T}}\mathbf{c} \right\},\end{aligned} \tag{14}$$

where

$$\hat{\mathbf{x}}(\mathbf{c}) = \max\big(\tilde{\mathbf{K}}_0^{\mathrm{T}}\mathbf{c}, 0\big). \tag{15}$$

The max is conducted elementwise in Eq. (15), so the operation zeros any negative elements in the $p \times 1$ input vector $\tilde{\mathbf{K}}_0^{\mathrm{T}}\mathbf{c}$. The minimization in the first line of Eq. (14) is referred to as the primal optimization problem in the primal variables $\mathbf{x}$. The second line of Eq. (14) is an equivalent optimization problem in both primal variables $\mathbf{x}$ ($p \times 1$) and dual variables $\mathbf{c}$ ($\tilde{p} \times 1$), the solution of which finds the saddle point ($\mathbf{x} = \hat{\mathbf{x}}$, $\mathbf{c} = \hat{\mathbf{c}}$) of its objective function. The inner maximization is achieved for given $\mathbf{x}$ at $\mathbf{c} = \hat{\mathbf{c}}(\mathbf{x})$, which recovers the primal problem from the saddle-point problem. The optimum $\tilde{p} \times 1$ dual vector is $\hat{\mathbf{c}}(\mathbf{x}) = (\tilde{\mathbf{g}} - \tilde{\mathbf{K}}_0\mathbf{x})/\lambda^2$ given the primal vector $\mathbf{x}$. This type of optimization transfer is a case of convex conjugation [46, Chap. 3], i.e., Fenchel-Legendre duality as noted in [47]. The saddle-point property permits interchange in the order of minimization and maximization [46, Chap. 5] in the second line of Eq. (14). The inner minimization is then over $\mathbf{x}$ for given $\mathbf{c}$: the minimizer is given by Eq. (15)

[27,47]. The maximization over the dual variables **c** in the third line of Eq. (14) is obtained from the saddle-point problem at this inner minimum. This maximization is an unconstrained dual optimization problem to the lower-bound-constrained primal optimization problem. The transformation of the primal problem into the dual problem in Eq. (14) is in essence the method of [27]. The optimal value of **c** ($\tilde{p} \times 1$) is determined by an unconstrained minimization [26] of the dual objective function in the third line of Eq. (14) multiplied by -1. The solution vector $\hat{\mathbf{x}}$ is obtained from this optimum $\hat{\mathbf{c}}$ via Eq. (15). The optimum λ^2 parameter can be determined by the method given in [26,27], or by the methods discussed in [44, Chap. 5]. We use generalized cross validation (GCV) [48,49] following the method in [28]. Application to 2D NMR data is straightforward (the matrix of measured 2D data is stacked as a vector, and the solution vector reshaped to a matrix) as indicated earlier in this section and detailed in [26].

An example of Tikhonov regularization applied to simulated 1D data containing contributions with positive and negative amplitude (see Fig. 1) is shown in Fig. 2. The non-negative and unconstrained versions of Tikhonov regularization are used here with λ^2 optimized by GCV in each case. The distribution in Fig. 2(b) (blue dashed line) was obtained by applying the method in [26,27], including a non-negativity constraint, to the simulated signal plus noise. The fitted distribution contained a single relaxation-time component (peak) of integral area 1.68 at a modal rate of $1/T = 1.3\ \mathrm{s}^{-1}$. The corresponding cumulative distribution plot is shown in Fig. 2(c) (blue dashed line). It is interesting to observe that the inclusion of the non-negativity constraint does not simply suppress the component of negative amplitude, but distorts the entire estimated distribution. The reason for this discrepancy is clear when the simulated and fitted data are compared in Fig. 1(a) with the fit residuals given in Fig. 2(a) (blue dashed line). The estimated data are a poor fit to the original signal with the result oversmoothed (large λ^2) to compensate for the subtractive component. This example is extreme, with the negative amplitude component having a significant magnitude compared to the total signal. Nevertheless, it highlights the limitation of this non-negative implementation of Tikhonov regularization. The distribution in Fig. 2(b) (red solid line) was calculated by Eq. (12). In this case, unconstrained Tikhonov regularization achieves a good fit to the original signal and the residuals in Fig. 2(a) (red solid line) are small, but the distribution is neither sparse nor sensible as it suffers from severe artificial ringing. The corresponding cumulative distribution shown in Fig. 2(c) (red solid line) also exhibits ringing, the full extent of which is clipped due to the axis scale. This standard method thus cannot re-mediate such reconstruction problems in general.

B. Half-bound Tikhonov regularization

If there is prior knowledge of the sign of f as it varies with T, then the above non-negative method can be adapted accordingly to yield "half-bound" Tikhonov regularization, which has distinct non-negative and nonpositive regions. Each column of $\tilde{\mathbf{K}}_0$ is multiplied by the assumed sign of $f(T)$ corresponding to the relevant value of T for that column, i.e., a mask with values $+1$ or -1 given by $\mathrm{sgn}(f(T))$, where sgn is the signum function. In this modified method, each element of **x** is thus proportional to the magnitude of $f(T)$ corresponding to the relevant value of T for that element. The non-negative Tikhonov reconstruction is run with the modified $\tilde{\mathbf{K}}_0$. The required result $\hat{f}(T)$ is recovered by multiplying each element of the estimated $\hat{\mathbf{x}}$ by $\mathrm{sgn}(f(T))$ corresponding to the relevant value of T for that element. The ill-posed nature of identifying $f(T)$ from Eq. (9) means that approximately equal magnitudes either side of a sign transition in $\hat{f}(T)$ correspond to canceling contributions with little net effect on the fitted signal. The corresponding kernel functions would be similar in magnitude and opposite in sign. Since the location of the sign transition is chosen, any resulting artifacts can be deleted manually near this known position. This change could of course result in the distortion or removal of nearby (genuine)

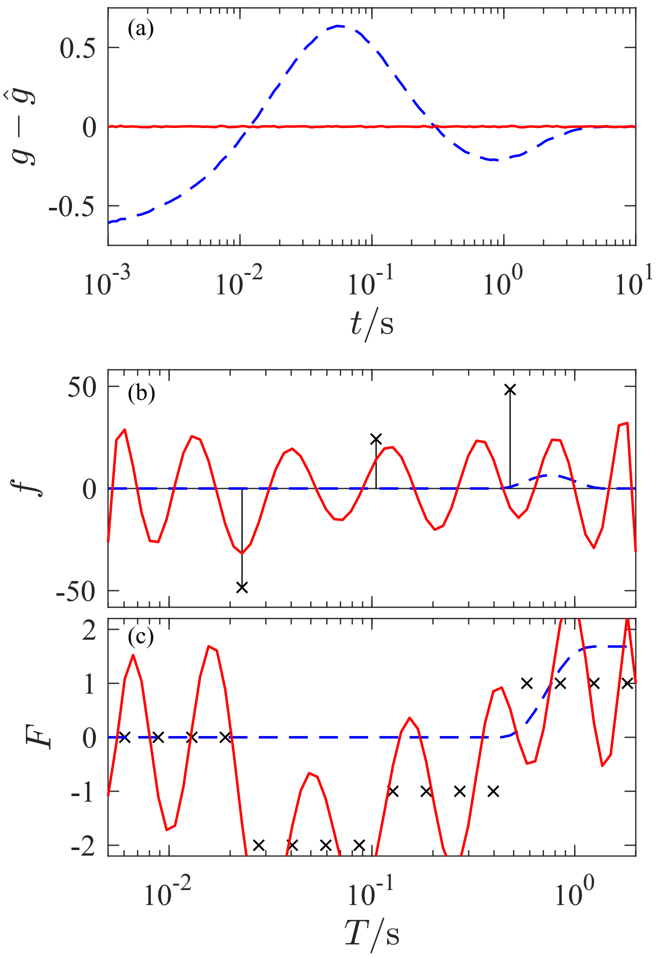

FIG. 2. Demonstration of Tikhonov regularization applied to the simulated 1D signal **g** shown in Fig. 1(a). Noise was added equivalent to a signal-to-noise ratio of SNR ≈ 1000. Estimated relaxation-time distributions $\hat{\mathbf{x}}$ are shown in (b) for non-negative Tikhonov regularization (blue dashed line) and unconstrained Tikhonov regularization (red solid line) with selected points from the ideal distribution (black × symbols) indicated for comparison. Vertical stems are included to guide the eye. The cumulative integral is shown for each of these distributions in (c) and the corresponding fit residuals $\mathbf{g} - \hat{\mathbf{g}}$ are shown in (a). The data $\hat{\mathbf{g}}$ obtained by non-negative Tikhonov regularization were included in Fig. 1(a). See online version for color.

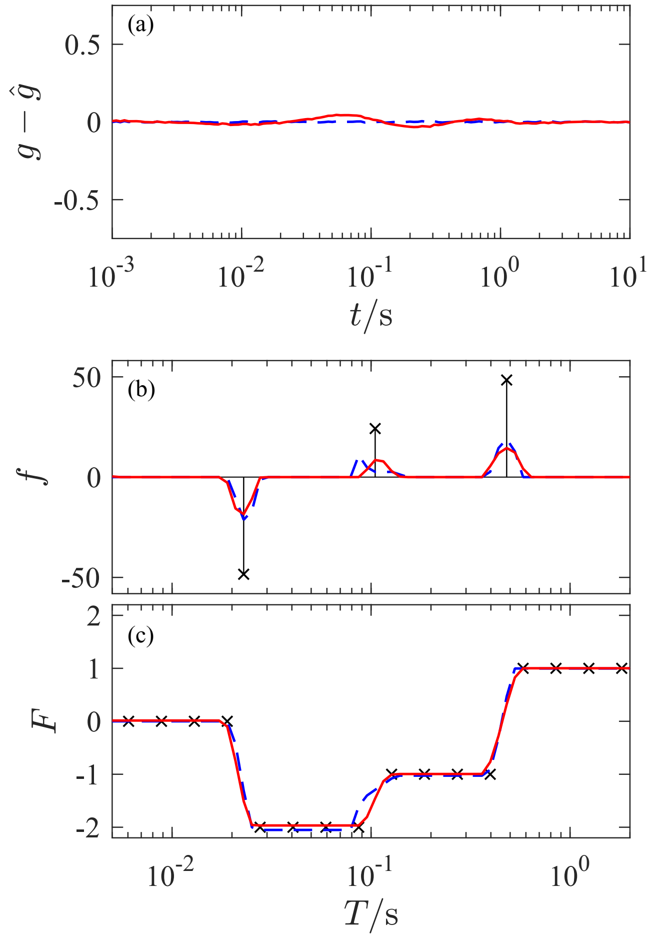

FIG. 3. Demonstration of half-bound Tikhonov regularization and primal-dual methods applied to the simulated 1D signal **g** shown in Fig. 1(a). Noise was added equivalent to a signal-to-noise ratio of SNR ≈ 1000. Estimated relaxation-time distributions $\hat{\mathbf{x}}$ are shown in (b) for half-bound Tikhonov regularization (blue dashed line) and the primal-dual algorithm (red solid line) with selected points from the ideal distribution (black × symbols) indicated for comparison. Vertical stems are included to guide the eye. The cumulative integral is shown for each of these distributions in (c) and the corresponding fit residuals $\mathbf{g} - \hat{\mathbf{g}}$ are shown in (a). See online version for color.

peaks. If required, $\hat{f}(T)$ can be rescaled to yield unit integral over $\log_{10} T$ after removal of the artifacts.

The half-bound Tikhonov method is applied to the aforementioned 1D simulated signal in Fig. 1, assuming nonpositive amplitudes for $T < 0.08$ s (non-negative elsewhere). The reconstructed distribution in Fig. 3(b) (blue dashed line) encompasses each of the required peaks, including the negative-amplitude component. The middle (positive) peak was slightly stretched towards the left (negative-amplitude) peak, but no spurious peaks appeared in this particular reconstruction. The peak integral areas (left to right) are −1.98, 1.03, 2.03, so the distortion does not impede quantification. The corresponding cumulative distribution shown in Fig. 3(c) (blue dashed line) reaches similar levels to the simulated distribution integral (indicated by the × symbols). Half-bound Tikhonov regularization thus reconstructs the simulated distribution including the negative-amplitude component. The fitted data closely follow the original signal as demonstrated by the residuals in Fig. 3(a) (blue dashed line) which are significantly better than those achieved by non-negative Tikhonov regularization [cf. Fig. 2(a)].

For the 2D case of a T_1-T_2 correlation in a diffusive-exchange system, components of the distribution are known to have negative amplitude [18] for apparent $T_2 > T_1$ if the two bulk relaxation times are similar. Such negative amplitudes in this setting are thus always associated with particular combinations of eigenvalues, specifically for the positions $\{T_{1,j}^{\mathrm{app}}, T_{2,i}^{\mathrm{app}}\}$ where $j > i$. This knowledge suggests that $\mathrm{sgn}(f(T_1, T_2)) = -1$ where $T_2 > T_1$ and $+1$ otherwise. However, the smoothing inherent in the Tikhonov regularization may result in positive-amplitude peaks near the $T_1 = T_2$ diagonal extending over this threshold. To avoid affecting the correct recovery of such peaks, the $\mathrm{sgn}(f(T_1, T_2))$ mask can instead be set to -1 over a smaller region of the distribution, say, where $T_2/T_1 > 0.5$, leaving the mask value as $+1$ elsewhere. The position of the sign transition can be adjusted for any given data set to minimize distortions in the estimated solution.

At the potential cost of some artifacts, the half-bound inversion retains the advantages of Tikhonov regularization: robust selection of the optimum smoothing parameter via GCV, minimal smoothing on the output, and sensitivity to low-amplitude contributions to the signal. Next, we apply a LASSO-type method to the same simulated signal for comparison.

C. Primal-dual algorithm

A generalization of the LASSO [50] encompasses the formulation

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \left\{ \|\mathbf{g} - \mathbf{K}_0\mathbf{x}\|_2^2 + \lambda_{\mathrm{I}} \|\mathbf{x}\|_1 + \lambda_{\mathrm{L}} \|\mathbf{L}_0\mathbf{x}\|_1 \right\}, \quad (16)$$

where λ_{I} and λ_{L} are both tuning parameters. Note the absence of the non-negativity constraint on $\mathbf{x}$ $(p \times 1)$. The bandwidth of $\log_{10} T$ is restricted in order to achieve a stable solution, as above with Tikhonov regularization. This problem becomes the standard LASSO of [32,33] for $\lambda_{\mathrm{L}} = 0$. Alternatively, it becomes anisotropic total-variation regularization if $\lambda_{\mathrm{I}} = 0$ and $\mathbf{L}_0$ is a first finite-difference matrix. A more sophisticated choice of $\mathbf{L}_0$, originally designed for imaging [51,52], was adopted in [40] to reconstruct distributions with both sparse and smooth contributions. For the purposes of avoiding spurious components, a second finite-difference matrix (with p columns) was used here, as is common in Tikhonov regularization to ensure smoothness [44, Chap. 8].

Primal-dual methods are applicable to a variety of optimization problems, including that for non-negative Tikhonov regularization in Eq. (14). For brevity, we hereafter use the term “primal-dual algorithm” specifically in the context of solving problems expressible in the form of Eq. (16), despite the generality of the term. The primal-dual algorithm of [38] was adapted in [39,40] to solve a problem of similar form to Eq. (16). The primal-dual algorithm in such a case is based on a saddle-point problem reformulated from the minimization in

Eq. (16) via convex conjugation, viz.,

$$\begin{aligned}\min_{\mathbf{x}}&\left\{\|\mathbf{g}-\mathbf{K}_0\mathbf{x}\|_2^2+\lambda_\mathrm{I}\|\mathbf{x}\|_1+\lambda_\mathrm{L}\|\mathbf{L}_0\mathbf{x}\|_1\right\}\\ &=\min_{\mathbf{x}}\max_{\mathbf{y}_\mathrm{I},\mathbf{y}_\mathrm{L}}\left\{\|\mathbf{g}-\mathbf{K}_0\mathbf{x}\|_2^2+\lambda_\mathrm{I}\left[\mathbf{y}_\mathrm{I}^\mathrm{T}\mathbf{x}-\Phi_1(\mathbf{y}_\mathrm{I})\right]\right.\\ &\quad\left.+\lambda_\mathrm{L}\left[\mathbf{y}_\mathrm{L}^\mathrm{T}\mathbf{L}_0\mathbf{x}-\Phi_1(\mathbf{y}_\mathrm{L})\right]\right\},\end{aligned}\tag{17}$$

where $\mathbf{y}_\mathrm{I}$ ($p\times 1$) and $\mathbf{y}_\mathrm{L}$ ($p\times 1$ assuming $\mathbf{L}_0$ is square) are vectors of dual variables. The indicator penalty function $\Phi_\alpha(\mathbf{y})$ yields ∞ if any element of $\mathbf{y}$ exceeds α in magnitude and yields zero otherwise (as a scalar result). The inner maximization in the second line of Eq. (17) is thus equivalent to maximizing $\lambda_\mathrm{I}\mathbf{y}_\mathrm{I}^\mathrm{T}\mathbf{x}$ subject to $\|\mathbf{y}_\mathrm{I}\|_\infty\leqslant 1$ and separately maximizing $\lambda_\mathrm{L}\mathbf{y}_\mathrm{L}^\mathrm{T}\mathbf{L}_0\mathbf{x}$ subject to $\|\mathbf{y}_\mathrm{L}\|_\infty\leqslant 1$, adding the results to the regression ℓ_2-norm term. The overall maximum is achieved at $\mathbf{y}_\mathrm{I}=\hat{\mathbf{y}}_\mathrm{I}(\mathbf{x})$ and $\mathbf{y}_\mathrm{L}=\hat{\mathbf{y}}_\mathrm{L}(\mathbf{x})$ for a given primal vector $\mathbf{x}$. The primal optimization in the first line of Eq. (17) is recovered from the saddle-point problem in the second line by noting $\hat{\mathbf{y}}_\mathrm{I}(\mathbf{x})=\mathrm{sgn}(\mathbf{x})$ and $\hat{\mathbf{y}}_\mathrm{L}(\mathbf{x})=\mathrm{sgn}(\mathbf{L}_0\mathbf{x})$ from the separate maximizations, given $\mathbf{x}$ (here $\hat{\mathbf{y}}_\mathrm{I}$ is $p\times 1$ and $\hat{\mathbf{y}}_\mathrm{L}$ is $p\times 1$ for square $\mathbf{L}_0$). The signum function is conducted elementwise on its input vector.

Optimization to find the saddle point in Eq. (17) is achieved by alternating updates of the primal variables ($\mathbf{x}$) and the dual variables ($\mathbf{y}_\mathrm{I}$ and $\mathbf{y}_\mathrm{L}$) [38]. The update of each dual vector is done using the proximity operation

$$\begin{aligned}\mathrm{prox}_1(\mathbf{y},\alpha)&=\arg\min_{\mathbf{u}}\left\{\tfrac{1}{2}\|\mathbf{u}-\mathbf{y}\|_2^2+\Phi_\alpha(\mathbf{u})\right\}\\ &=\mathrm{sgn}(\mathbf{y})\odot\min(\mathrm{sgn}(\mathbf{y})\odot\mathbf{y},\alpha),\end{aligned}\tag{18}$$

where α is the threshold parameter. This operation is a projection of the vector $\mathbf{y}$ on to the convex set of vectors with ℓ_∞ norm, $\|\mathbf{y}\|_\infty\leqslant\alpha$, and is a form of hard threshold on the element magnitudes. Both sgn and min are conducted elementwise in Eq. (18), so the $\mathrm{prox}_1(\mathbf{y})$ operation separately prunes each element of $\mathbf{y}$ to absolute value α at most, viz., $\mathrm{sgn}(y_i)\min(|y_i|,\alpha)$. The update of the primal vector uses the proximity operation

$$\begin{aligned}\mathrm{prox}_2(\mathbf{x},\beta)&=\arg\min_{\mathbf{z}}\left\{\tfrac{1}{2}\|\mathbf{z}-\mathbf{x}\|_2^2+\beta\|\mathbf{g}-\mathbf{K}_0\mathbf{z}\|_2^2\right\}\\ &=\mathbf{x}+\mathbf{V}_0\big(\mathbf{w}_0\odot\big(\mathbf{U}_0^\mathrm{T}(\mathbf{g}-\mathbf{K}_0\mathbf{x})\big)\big),\end{aligned}\tag{19}$$

where the elements of $p\times 1$ vector $\mathbf{w}_0$ are $w_{0,i}=s_{0,i}/[s_{0,i}^2+1/(2\beta)]$ and β is a parameter that determines the amount of regularization. This operation is a version of unconstrained Tikhonov regularization with tuning parameter $\lambda^2=1/(2\beta)$, so it is straightforward to solve using the SVD of $\mathbf{K}_0$.

The iterative minimization is initialized with $p\times 1$ zero vectors for $\mathbf{x}^{(0)}$, and $\mathbf{y}_\mathrm{I}^{(0)}$ and $\mathbf{y}_\mathrm{L}^{(0)}$. The lth estimates of the dual vectors $\mathbf{y}_\mathrm{I}^{(l)}$ and $\mathbf{y}_\mathrm{L}^{(l)}$ are calculated by

$$\mathbf{y}_\mathrm{I}^{(l)}=\mathrm{prox}_1\big(\mathbf{y}_\mathrm{I}^{(l-1)}+\sigma\lambda_\mathrm{I}\mathbf{z}^{(l-1)},1\big),\tag{20}$$

$$\mathbf{y}_\mathrm{L}^{(l)}=\mathrm{prox}_1\big(\mathbf{y}_\mathrm{L}^{(l-1)}+\sigma\lambda_\mathrm{L}\mathbf{L}_0\mathbf{z}^{(l-1)},1\big),\tag{21}$$

where $\sigma=1/(\Lambda^2\tau)$ with $\Lambda=\max(\lambda_\mathrm{I},\lambda_\mathrm{L}|\mathrm{maxeig}(\mathbf{L}_0)|)$, $\tau=10^5/(2\,\mathrm{maxeig}(\mathbf{K}_0^\mathrm{T}\mathbf{K}_0))$, and maxeig is the maximum magnitude eigenvalue of its matrix argument. The distribution estimate is updated as

$$\mathbf{x}^{(l)}=\mathrm{prox}_2\big(\mathbf{x}^{(l-1)}-\tau\big[\lambda_\mathrm{I}\mathbf{y}_\mathrm{I}^{(l)}+\lambda_\mathrm{L}\mathbf{L}_0^\mathrm{T}\mathbf{y}_\mathrm{L}^{(l)}\big],\tau\big),\tag{22}$$

$$\mathbf{z}^{(l)}=\mathbf{x}^{(l)}+\theta(\mathbf{x}^{(l)}-\mathbf{x}^{(l-1)}).\tag{23}$$

While $\mathbf{x}^{(l)}$ ($p\times 1$) contains the desired estimate of the distribution f within scaling, the iteration is accelerated by calculating an auxiliary version of distribution $\mathbf{z}^{(l)}$ ($p\times 1$) via over-relaxation in Eq. (23). The next iteration proceeds with this vector of auxiliary variables as input. A possible choice for the over-relaxation parameter is $\theta=1/\sqrt{1+2\gamma\tau}$ [38] where the dual Lipschitz constant is $\gamma=2\;\mathrm{mineig}(\mathbf{K}_0^\mathrm{T}\mathbf{K}_0)/\Lambda$, and mineig is the minimum magnitude eigenvalue of its matrix argument. For inversions in NMR relaxation experiments, the matrix $\mathbf{K}_0^\mathrm{T}\mathbf{K}_0$ is often sufficiently ill conditioned (maxeig $\gg$ mineig) that the calculation yields $\theta\approx 1$. The SVD of $\mathbf{K}_0$ is again useful for these calculations since $\mathrm{maxeig}(\mathbf{K}_0^\mathrm{T}\mathbf{K}_0)=s_{0,\max}^2$ and $\mathrm{mineig}(\mathbf{K}_0^\mathrm{T}\mathbf{K}_0)=s_{0,\min}^2$.

The above algorithm typically needs to be rerun over various values of λ_I because a suitable value is not typically known in advance. A limited number of iterations, say $l_{\max}=10^3$, are performed at each λ_I value, and the resulting $\mathbf{x}^{(l_{\max})}$ and $\mathbf{y}_\mathrm{I}^{(l_{\max})}$ and $\mathbf{y}_\mathrm{L}^{(l_{\max})}$ are used instead of zero vectors to initialize the minimization for the next λ_I. A reasonable starting estimate λ_A for the penalty parameter is the estimated root-mean-square (rms) noise in measurements $\mathbf{g}$ [53]. The optimum solution is found by incrementing λ_I from λ_A to $\lambda_\mathrm{B}=2\lambda_\mathrm{A}s_{0,\max}$ over logarithmic intervals. The optimum parameter $\lambda_{\mathrm{I,opt}}$ is chosen where a sudden increase occurs in a plot of $\log\|\mathbf{x}(\lambda_\mathrm{I})\|_1$ versus $\log\|\mathbf{g}-\mathbf{K}_0\mathbf{x}(\lambda_\mathrm{I})\|_2$. This grid search is a form of L-curve method [44, Chap. 5] for tuning λ_I. Instead of conducting a separate grid search, the value of λ_L can, for example, be chosen in fixed proportion to λ_I [40].

An example inversion of 1D data using the primal-dual algorithm is given in Fig. 3. For this 1D case, λ_L was set to zero so the corresponding step in Eq. (21) was ignored. The primal-dual algorithm recovers the approximate peak positions [see the distribution in Fig. 3(b) (red solid line)]. The distribution peaks are correctly centered on the original rate values. The peak integral areas (left to right) are −1.97, 0.97, 2.00 and the corresponding cumulative distribution in Fig. 3(c) (red solid line) reaches similar levels to the broadened version of the original simulated distribution (indicated by the × symbols). The primal-dual algorithm thus recovers the approximate component amplitudes. The accuracy of these amplitudes is similar to the half-bound Tikhonov result. The fitted data are a good approximation to the original signal with small fit residuals [see Fig. 3(a) (red solid line)]. The primal-dual algorithm reconstructs the simulated distribution including the negative-amplitude component.

The primal-dual algorithm could be implemented on 2D NMR data by stacking the $q_2\times q_1$ data matrix $\mathbf{G}$ columnwise as a $q\times 1$ vector ($q=q_1q_2$). The resulting $p\times 1$ solution vector can be reshaped into a $p_2\times p_1$ distribution matrix $\mathbf{X}$ ($p=p_1p_2$). When the kernel function is separable, as in the T_1-T_2 correlation, it is computationally efficient to solve the generalized LASSO in the equivalent matrix form

$$\min_{\mathbf{X}}\Big\{\big\|\mathbf{G}-\mathbf{K}_2\mathbf{X}\mathbf{K}_1^{\mathrm{T}}\big\|_{\mathrm{F}}^2+\lambda_{\mathrm{I}}\|\mathrm{vec}(\mathbf{X})\|_1 + \lambda_{\mathrm{L}}\big[\big\|\mathrm{vec}\big(\mathbf{X}\mathbf{L}_1^{\mathrm{T}}\big)\big\|_1+\|\mathrm{vec}(\mathbf{L}_2\mathbf{X})\|_1\big]\Big\}, \tag{24}$$

where $\mathbf{K}_1$ is the $q_1 \times p_1$ kernel matrix for the indirect dimension (T_1) and $\mathbf{K}_2$ is the $q_2 \times p_2$ kernel matrix for the direct dimension (T_2). Both $\mathbf{L}_1$ ($p_1 \times p_1$) and $\mathbf{L}_2$ ($p_2 \times p_2$) are second finite-difference matrices, which contribute to the 2D finite-difference Laplacian operator $\mathbf{L}_0 = \mathbf{I}_1 \otimes \mathbf{L}_2 + \mathbf{L}_1 \otimes \mathbf{I}_2$ ($p \times p$). The identity matrices $\mathbf{I}_1$ and $\mathbf{I}_2$ have respective sizes $p_1 \times p_1$ and $p_2 \times p_2$ here, but there is no need to form $\mathbf{L}_0$ directly for these calculations. Similar to Eq. (18), prox_1 is conducted elementwise for matrix inputs. The matrix-input version of prox_2 is

$$\begin{aligned}\mathrm{prox}_2(\mathbf{X},\beta) &= \arg\min_{\mathbf{Z}}\big\{\tfrac{1}{2}\|\mathbf{Z}-\mathbf{X}\|_{\mathrm{F}}^2+\beta\big\|\mathbf{G}-\mathbf{K}_2\mathbf{Z}\mathbf{K}_1^{\mathrm{T}}\big\|_{\mathrm{F}}^2\big\}\\ &= \mathbf{X}+\mathbf{V}_2\big(\mathbf{W}_0\odot\big(\mathbf{U}_2^{\mathrm{T}}(\mathbf{G}-\mathbf{K}_2\mathbf{X}\mathbf{K}_1^{\mathrm{T}})\mathbf{U}_1\big)\big)\mathbf{V}_1^{\mathrm{T}},\end{aligned} \tag{25}$$

where the elements of $p \times p$ matrix $\mathbf{W}_0$ are $W_{0,ij} = s_{2,i}s_{1,j}/[s_{2,i}^2 s_{1,j}^2 + 1/(2\beta)]$. The primal-dual algorithm is implemented with $p_2 \times p_1$ matrices $\mathbf{X}^{(l)}$, $\mathbf{Z}^{(l)}$, $\mathbf{Y}_{\mathrm{I}}^{(l)}$ and $\mathbf{Y}_{\mathrm{L}}^{(l)}$ (all zero for $l = 0$ initially) as

$$\mathbf{Y}_{\mathrm{I}}^{(l)} = \mathrm{prox}_1\big(\mathbf{Y}_{\mathrm{I}}^{(l-1)}+\sigma\lambda_{\mathrm{I}}\mathbf{Z}^{(l-1)},1\big), \tag{26}$$

$$\mathbf{Y}_{\mathrm{L}}^{(l)} = \mathrm{prox}_1\big(\mathbf{Y}_{\mathrm{L}}^{(l-1)}+\sigma\lambda_{\mathrm{L}}\big(\mathbf{L}_2\mathbf{Z}^{(l-1)}+\mathbf{Z}^{(l-1)}\mathbf{L}_1^{\mathrm{T}}\big),1\big), \tag{27}$$

$$\mathbf{X}^{(l)} = \mathrm{prox}_2\big(\mathbf{X}^{(l-1)}-\tau\big[\lambda_{\mathrm{I}}\mathbf{Y}_{\mathrm{I}}^{(l)}\lambda_{\mathrm{L}}\big(\mathbf{L}_2^{\mathrm{T}}\mathbf{Y}_{\mathrm{L}}^{(l)}+\mathbf{Y}_{\mathrm{L}}^{(l)}\mathbf{L}_1\big)\big],\tau\big), \tag{28}$$

$$\mathbf{Z}^{(l)} = \mathbf{X}^{(l)}+\theta(\mathbf{X}^{(l)}-\mathbf{X}^{(l-1)}), \tag{29}$$

where θ is set equal to unity. We use this efficient matrix-based approach to estimate optimum solutions for T_1-T_2 correlation data, setting $\lambda_{\mathrm{L}} = \lambda_{\mathrm{I}}/10$ in the grid search for λ_{I}.

III. EXCHANGE MODELS: SIMULATIONS AND ANALYTIC SOLUTIONS

In simple geometries, the eigenfunctions and eigenvalues can be calculated explicitly. However, for more complicated scenarios, generation of time-dependent magnetization is achieved more readily through simulation. We turn to the finite-element multiphysics modeling package COMSOL as a convenient method for simulating spin relaxation in nontrivial situations, in particular coupled pores in the slow diffusion regime. A built-in partial differential equation (PDE) solver was used to calculate the time-dependent evolution of the magnetization according the Bloch-Torrey equation. The longitudinal magnetization was modeled as $m(\mathbf{r},t) \propto M_z$ in Eq. (1) and surface relaxation imposed by Eq. (2) with a surface relaxivity ρ_1. Similarly, the transverse magnetization (always real valued in these simulations) was modeled as $m(\mathbf{r},t) \propto M_x$ (on resonance in the rotating frame), with surface relaxivity ρ_2, to satisfy Eqs. (3) and (4). The scaling of the maximum coherent magnetization was defined exactly with $\int m(\mathbf{r},0)dV(\mathbf{r}) = 1$ in all simulations.

The T_1-T_2 experiment was simulated using two separate COMSOLsolvers, one for each stage of relaxation. The longitudinal magnetization was first calculated as a function of recovery time t_1, incremented logarithmically from 10^{-4} s to 10 s in 128 steps. The value of M_z at each t_1 time was then used as the input $M_z \rightarrow M_x$ for the calculation of the transverse magnetization decays with data stored at times t_2 incremented logarithmically from 10^{-4} s to 10 s in 128 steps. To simplify the generation of 2D data, the COMSOL server was driven from a MATLABscript. The basic geometries used in the simulations are illustrated in Fig. 4. The same bulk relaxation rate was assigned to longitudinal and transverse processes for the fluid in the pore volume in all cases: $T_1^{\mathrm{b}} = T_2^{\mathrm{b}} = 2$ s.

A. Isolated spherical pore

Analytic solutions to the Bloch-Torrey equation are available for a spherical geometry [24,54,55]. The eigenfunctions of the Laplacian operator (∇^2) for a given boundary condition in a spherical geometry are products of spherical Bessel functions of the first kind and spherical harmonics; the corresponding eigenvalues are nonpositive. These eigenfunctions are complex valued in general and are orthogonal with respect to integration over the volume of the ball $V_{\mathrm{sph}} = 4\pi r_{\mathrm{sph}}^3/3$. The rotational symmetry here means the only eigenfunctions that contribute are those proportional to the zeroth-order spherical Bessel function of the first kind, $\mathrm{j}_0(x) = \sin(x)/x$, which is identical to the sinc function. The contribution to the Fickian diffusive flux, from each of these eigenfunctions, is proportional to the first-order spherical Bessel function of the first kind $\mathrm{j}_1(x) = -\mathrm{j}_0'(x) = \sin(x)/x^2 - \cos(x)/x$. The spatial dependence is only with the radial coordinate for these eigenfunctions, viz., $\mathrm{j}_0(\xi r/r_{\mathrm{sph}})/4\pi$. The contributing eigenfunctions in this rotationally symmetric situation are thus real valued. The allowable values of ξ are those that satisfy the Robin boundary condition at $r = r_{\mathrm{sph}}$, which is $-\mathcal{D}\xi\mathrm{j}_0'(\xi)/r_{\mathrm{sph}} - \rho\mathrm{j}_0(\xi) = 0$ over the entire spherical surface with area $S_{\mathrm{sph}} = 4\pi r_{\mathrm{sph}}^2$. These discrete values of ξ are thus the positive roots of

$$\frac{\rho r_{\mathrm{sph}}}{\mathcal{D}} = \frac{\xi\mathrm{j}_1(\xi)}{\mathrm{j}_0(\xi)} = 1-\xi\cot(\xi). \tag{30}$$

The corresponding eigenvalues of the diffusion operator ($-\mathcal{D}\nabla^2$) with the above boundary condition are determined by these roots as $\mathcal{D}\xi^2/r_{\mathrm{sph}}^2$. If $\rho r_{\mathrm{sph}}/\mathcal{D} \ll 1$, the lowest eigenvalue is approximately $\rho S_{\mathrm{sph}}/V_{\mathrm{sph}} = 3\rho/r_{\mathrm{sph}}$ and becomes zero in the limit $\rho \rightarrow 0$ (viz. the Neumann boundary condition). To provide the appropriate normalization of the eigenfunctions over the volume of the sphere, it is convenient to define

$$\eta(\xi,\delta) = \frac{2}{3}\frac{\xi^2}{\xi^2+\delta(\delta-1)}. \tag{31}$$

The two sets of eigenfunctions for longitudinal and transverse relaxation have similar functional form, but differ due to the different sets of ξ values involved, depending on the surface relaxivity. In each case, the roots are indexed in increasing order with $j = 0$ for the smallest root for longitudinal relaxation and $i = 0$ for the smallest root for transverse relaxation. The two sets of eigenvalues are interleaved:

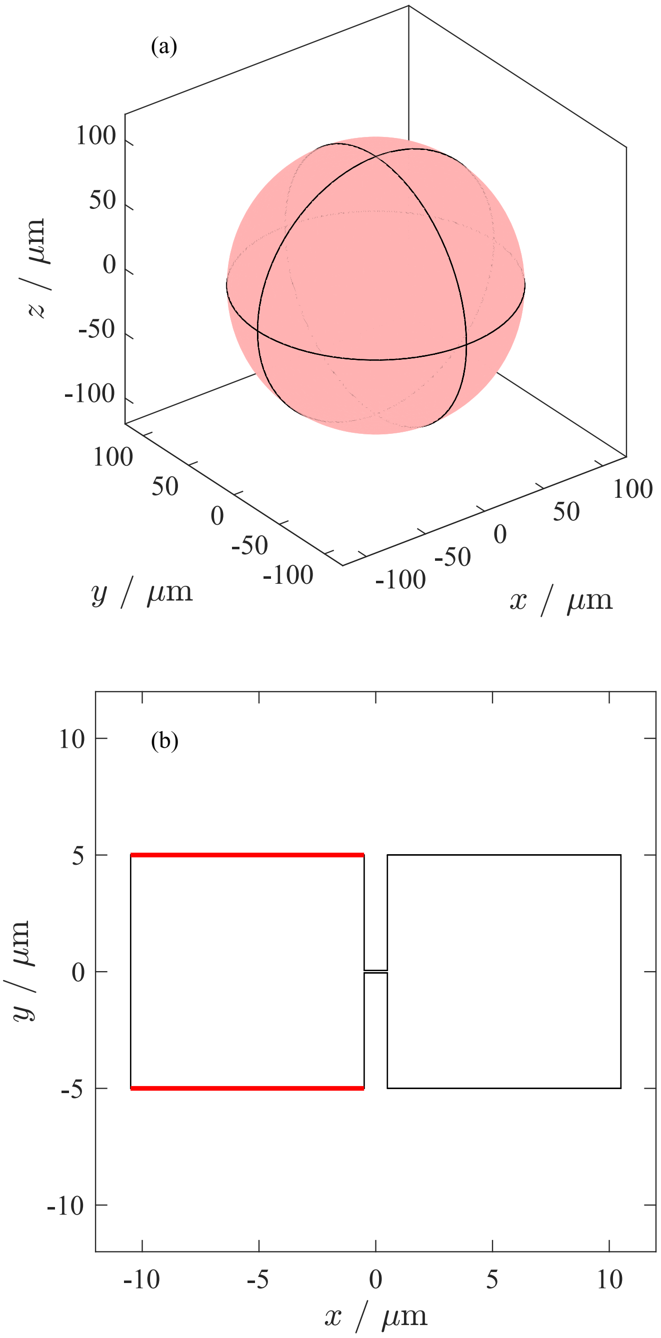

FIG. 4. COMSOL simulation geometries for (a) 3D spherical pore and (b) 2D weakly coupled pores. The relaxing boundaries are highlighted red in each geometry. See online version for color.

$\xi_{1,j} < \xi_{2,j} < \xi_{1,j+1} < \xi_{2,j+1}$, given $\rho_1 < \rho_2$. For diffusion with longitudinal relaxation, each root ($\xi_{1,j}$) of Eq. (30) with $\rho = \rho_1$ yields an eigenfunction

$$\psi_j(r) = \sqrt{\frac{\eta\left(\xi_{1,j}, \frac{\rho_1 r_{\rm sph}}{\mathcal{D}}\right)}{V_{\rm sph}}}\frac{{\rm j}_0\left(\xi_{1,j}\frac{r}{r_{\rm sph}}\right)}{{\rm j}_0(\xi_{1,j})}, \tag{32}$$

with $\int_{r=0}^{r_{\rm sph}} |\psi_j(r)|^2 4\pi r^2 dr = 1$ and corresponding eigenvalue

$$\frac{1}{T_{1,j}^{\rm app}} = \frac{\mathcal{D}\xi_{1,j}^2}{r_{\rm sph}^2} + \frac{1}{T_1^{\rm b}}. \tag{33}$$

For diffusion with transverse relaxation, each root ($\xi_{2,i}$) of Eq. (30) with $\rho = \rho_2$ yields an eigenfunction

$$\phi_i(r) = \sqrt{\frac{\eta\left(\xi_{2,i}, \frac{\rho_2 r_{\rm sph}}{\mathcal{D}}\right)}{V_{\rm sph}}}\frac{{\rm j}_0\left(\xi_{2,i}\frac{r}{r_{\rm sph}}\right)}{{\rm j}_0(\xi_{2,i})}, \tag{34}$$

with $\int_{r=0}^{r_{\rm sph}} |\phi_i(r)|^2 4\pi r^2 dr = 1$ and corresponding eigenvalue

$$\frac{1}{T_{2,i}^{\rm app}} = \frac{\mathcal{D}\xi_{2,i}^2}{r_{\rm sph}^2} + \frac{1}{T_2^{\rm b}}. \tag{35}$$

The analytic amplitude for each peak in the T_1-T_2 correlation with position $\{T_{1,j}^{\rm app}, T_{2,i}^{\rm app}\}$ is $A_{ij} = d_i Q_{ij} a_j$, where for $\rho_1, \rho_2 > 0$ the required integrals are

$$\begin{aligned} a_j &= \int_{r=0}^{r_{\rm sph}} \psi_j^*(r)\frac{1}{V_{\rm sph}} 4\pi r^2 dr = \sqrt{\frac{\eta\left(\xi_{1,j}, \frac{\rho_1 r_{\rm sph}}{\mathcal{D}}\right)}{V_{\rm sph}}}\frac{3\,{\rm j}_1(\xi_{1,j})}{\xi_{1,j}\,{\rm j}_0(\xi_{1,j})} \\ &= \sqrt{\frac{\eta\left(\xi_{1,j}, \frac{\rho_1 r_{\rm sph}}{\mathcal{D}}\right)}{V_{\rm sph}}}\frac{3}{\xi_{1,j}^2}\frac{\rho_1 r_{\rm sph}}{\mathcal{D}}, \end{aligned} \tag{36}$$

$$\begin{aligned} d_i &= \int_{r=0}^{r_{\rm sph}} \phi_i(r) 4\pi r^2 dr = \sqrt{V_{\rm sph}\eta\left(\xi_{2,i}, \frac{\rho_2 r_{\rm sph}}{\mathcal{D}}\right)}\frac{3\,{\rm j}_1(\xi_{2,i})}{\xi_{2,i}\,{\rm j}_0(\xi_{2,i})} \\ &= \sqrt{V_{\rm sph}\eta\left(\xi_{2,i}, \frac{\rho_2 r_{\rm sph}}{\mathcal{D}}\right)}\frac{3}{\xi_{2,i}^2}\frac{\rho_2 r_{\rm sph}}{\mathcal{D}}, \end{aligned} \tag{37}$$

and assuming $\rho_1 \neq \rho_2$,

$$\begin{aligned} Q_{ij} &= \int_{r=0}^{r_{\rm sph}} \phi_j^*(r)\psi_j(r) 4\pi r^2 dr \\ &= \sqrt{\eta\left(\xi_{2,i}, \frac{\rho_2 r_{\rm sph}}{\mathcal{D}}\right)\eta\left(\xi_{1,j}, \frac{\rho_1 r_{\rm sph}}{\mathcal{D}}\right)} \\ &\quad \times \left(\frac{\xi_{2,i}^2}{\xi_{2,i}^2 - \xi_{1,j}^2}\frac{3\,{\rm j}_1(\xi_{2,i})}{\xi_{2,i}\,{\rm j}_0(\xi_{2,i})} - \frac{\xi_{1,j}^2}{\xi_{2,i}^2 - \xi_{1,j}^2}\frac{3\,{\rm j}_1(\xi_{1,j})}{\xi_{1,j}\,{\rm j}_0(\xi_{1,j})}\right) \\ &= \sqrt{\eta\left(\xi_{2,i}, \frac{\rho_2 r_{\rm sph}}{\mathcal{D}}\right)\eta\left(\xi_{1,j}, \frac{\rho_1 r_{\rm sph}}{\mathcal{D}}\right)}\frac{3\frac{\rho_2 r_{\rm sph}}{\mathcal{D}} - 3\frac{\rho_1 r_{\rm sph}}{\mathcal{D}}}{\xi_{2,i}^2 - \xi_{1,j}^2}. \end{aligned} \tag{38}$$

Note that when $\xi_{1,j} > \xi_{2,i}$, $Q_{ij} < 0$ (given $\rho_1 < \rho_2$) and if, in addition, $T_1^{\rm b} = T_2^{\rm b}$, then $T_{1,j}^{\rm app} < T_{2,i}^{\rm app}$. The analytic peak amplitudes for the T_1-T_2 correlation were calculated using the nine lowest modes ($i, j = 0, 1, 2, \ldots, 8$) of the eigenfunctions. The diffusion regimes are dictated by the ratio $\rho r_{\rm sph}/\mathcal{D}$; here, $r_{\rm sph}$ and $\mathcal{D}$ are constant. For example, we represent fast diffusion by $\rho_2 r_{\rm sph}/\mathcal{D} = 0.25$, intermediate diffusion by $\rho_2 r_{\rm sph}/\mathcal{D} = 2.5$, and slow diffusion by $\rho_2 r_{\rm sph}/\mathcal{D} = 25$.

To simulate time-domain data corresponding to the Brownstein-Tarr diffusion regimes, a 3D spherical pore of fixed radius $r_{\rm sph} = 100\ \mu$m was used [see Fig. 4(a)]. The sphere was taken to be filled with liquid with $\mathcal{D} = 2 \times 10^{-9}$ m^2 s^{-1} and the Bloch-Torrey equations were solved within the geometry. The surface of the sphere was assigned relaxivities of $\rho_1 = 1\ \mu$m s^{-1} (fast diffusion), $\rho_1 = 10\ \mu$m s^{-1} (intermediate diffusion), or $\rho_1 = 100\ \mu$m s^{-1} (slow diffusion); $\rho_2 = 5\rho_1$ always. The time-domain signal data were used as input for inversion to estimate the T_1-T_2 correlation distribution. No noise was added to these simulated data in order to identify any inherent inaccuracies or discrepancies between the simulation and analytic predictions. Inversion of these simulated data was performed separately by non-negative Tikhonov regularization, half-bound Tikhonov regularization, and the primal-dual method.

B. Two coupled pores

In this simple model, two pores are connected by a thin throat of length L_t. Pore A has volume V_A and pore B has volume V_B. The boundary of pore A has total surface area S_A and the boundary of pore B has total surface area S_B. Exchange from pore A to B is described by the rate K_{AB}, and exchange from B to A by K_{BA}, such that $K_{AB}V_A = K_{BA}V_B$ for detailed balance. These rates are controlled by the mass transfer coefficient $k_m = \mathcal{D}/L_t$ and the pore throat cross-sectional area S_{AB}^{cross} which throttles the exchange: $K_{AB} = k_m S_{AB}^{cross}/V_A$ and $K_{BA} = k_m S_{AB}^{cross}/V_B$. Surface relaxation in pore A takes place on an active surface area S_A^{relax}. Surface relaxation in pore B takes place on an active surface area S_B^{relax}. Each pore is assumed to be in the fast diffusion regime, which requires $\rho_1 S_A^{relax} V_A/(\mathcal{D}S_A^2)$, $\rho_1 S_B^{relax} V_B/(\mathcal{D}S_B^2) \ll 1$, $\rho_2 S_A^{relax} V_A/(\mathcal{D}S_A^2)$, $\rho_2 S_B^{relax} V_B/(\mathcal{D}S_B^2) \ll 1$, and $k_m S_{AB}^{cross} V_A/(\mathcal{D}S_A^2)$, $k_m S_{AB}^{cross} V_B/(\mathcal{D}S_B^2) \ll 1$. Diffusion within each pore is thus sufficiently fast that the magnetizations are uniform within each pore, albeit the two pores have different magnetizations in general. The time-dependent scaled magnetizations for the two pores are $m_{1,A}(t)$ and $m_{1,B}(t)$ for the longitudinal relaxation part of the T_1-T_2 experiment, with $m_{2,A}(t)$ and $m_{2,B}(t)$ used for the transverse relaxation part. The initial scaled magnetizations in each pore are the same: $m_{1,A}(0) = m_{1,B}(0) = -1/(V_A + V_B)$ following perfect inversion of the equilibrium magnetization $m^{eq} = 1/(V_A + V_B)$. In the absence of interpore exchange (i.e., decoupled pores, $k_m \to 0$), the isolated-pore longitudinal relaxation rates are simply $R_{1,A} = \rho_1 S_A^{relax}/V_A + 1/T_1^b$ and $R_{1,B} = \rho_1 S_B^{relax}/V_B + 1/T_1^b$, and the isolated-pore transverse relaxation rates are $R_{2,A} = \rho_2 S_A^{relax}/V_A + 1/T_2^b$ and $R_{2,B} = \rho_2 S_B^{relax}/V_B + 1/T_2^b$. The apparent relaxation decays, in the presence of exchange between the weakly coupled pores, are obtained by solving the systems of ordinary differential equations (ODEs) for both longitudinal and transverse cases. For inversion recovery

$$\frac{d}{dt}\begin{pmatrix} m_{1,A}(t) - m^{eq} \\ m_{1,B}(t) - m^{eq} \end{pmatrix} = \mathbf{C}_1 \begin{pmatrix} m_{1,A}(t) - m^{eq} \\ m_{1,B}(t) - m^{eq} \end{pmatrix}, \tag{39}$$

where the relaxation-coupling matrix is

$$\mathbf{C}_1 = \begin{pmatrix} -R_{1,A} - K_{AB} & K_{AB} \\ K_{BA} & -R_{1,B} - K_{BA} \end{pmatrix}. \tag{40}$$

A corresponding system of ODEs exists for transverse relaxation:

$$\frac{d}{dt}\begin{pmatrix} m_{2,A}(t) \\ m_{2,B}(t) \end{pmatrix} = \mathbf{C}_2 \begin{pmatrix} m_{2,A}(t) \\ m_{2,B}(t) \end{pmatrix}, \tag{41}$$

where

$$\mathbf{C}_2 = \begin{pmatrix} -R_{2,A} - K_{AB} & K_{AB} \\ K_{BA} & -R_{2,B} - K_{BA} \end{pmatrix}. \tag{42}$$

The relaxation rates in the T_1-T_2 correlation [6,29] correspond to the eigenvalues of $-\mathbf{C}_1$ and $-\mathbf{C}_2$, respectively, $1/T_{1,j}^{app}$ and $1/T_{2,i}^{app}$, noting that indices i, j only take the values 0 or 1 in this situation. Expressions for the peak amplitudes are also given in [6,29], which can be evaluated from the eigenvectors of $-\mathbf{C}_1$ and $-\mathbf{C}_2$. The components of the eigenvector $(\psi_{0,A}, \psi_{0,B})^T$ of $-\mathbf{C}_1$ are given by

$$\psi_{0,A} = \sqrt{\frac{R_{1,B} + K_{BA} - \frac{1}{T_{1,0}^{app}}}{\frac{V_A}{T_{1,1}^{app}} - \frac{V_A}{T_{1,0}^{app}}}}, \tag{43}$$

$$\psi_{0,B} = \sqrt{\frac{R_{1,A} + K_{AB} - \frac{1}{T_{1,0}^{app}}}{\frac{V_B}{T_{1,1}^{app}} - \frac{V_B}{T_{1,0}^{app}}}}, \tag{44}$$

with $V_A|\psi_{0,A}|^2 + V_B|\psi_{0,B}|^2 = 1$, and the corresponding eigenvalue is

$$\frac{1}{T_{1,0}^{app}} = \frac{1}{2}(R_{1,A} + K_{AB} + R_{1,B} + K_{BA}) - \frac{1}{2}\sqrt{(R_{1,A} + K_{AB} - R_{1,B} - K_{BA})^2 + 4K_{AB}K_{BA}}. \tag{45}$$

The components of the eigenvector $(\psi_{1,A}, \psi_{1,B})^T$ of $-\mathbf{C}_1$ are given by

$$\psi_{1,A} = -\sqrt{\frac{R_{1,A} + K_{AB} - \frac{1}{T_{1,0}^{app}}}{\frac{V_A}{T_{1,1}^{app}} - \frac{V_A}{T_{1,0}^{app}}}}, \tag{46}$$

$$\psi_{1,B} = \sqrt{\frac{R_{1,B} + K_{BA} - \frac{1}{T_{1,0}^{app}}}{\frac{V_B}{T_{1,1}^{app}} - \frac{V_B}{T_{1,0}^{app}}}}, \tag{47}$$

with $V_A|\psi_{1,A}|^2 + V_B|\psi_{1,B}|^2 = 1$, and the corresponding eigenvalue is

$$\frac{1}{T_{1,1}^{app}} = \frac{1}{2}(R_{1,A} + K_{AB} + R_{1,B} + K_{BA}) + \frac{1}{2}\sqrt{(R_{1,A} + K_{AB} - R_{1,B} - K_{BA})^2 + 4K_{AB}K_{BA}}. \tag{48}$$

These eigenvectors also satisfy $\psi_{0,A}^* V_A \psi_{1,A} + \psi_{0,B}^* V_B \psi_{1,B} = 0$. The system of ODEs in Eq. (39) can thus be solved analytically to give

$$\begin{aligned} \begin{pmatrix} m_{1,A}(t_1) \\ m_{1,B}(t_1) \end{pmatrix} &= \begin{pmatrix} m^{eq} \\ m^{eq} \end{pmatrix} + \exp\{\mathbf{C}_1 t_1\} \begin{pmatrix} m_{1,A}(0) - m^{eq} \\ m_{1,B}(0) - m^{eq} \end{pmatrix} \\ &= \sum_{j=0}^{1} \begin{pmatrix} \psi_{j,A} \\ \psi_{j,B} \end{pmatrix} \left[1 - 2\exp\left\{-\frac{t_1}{T_{1,j}^{app}}\right\}\right] a_j, \end{aligned} \tag{49}$$

where $a_j = \psi_{j,A}^* V_A m^{eq} + \psi_{j,B}^* V_B m^{eq}$. The components of the eigenvector $(\phi_{0,A}, \phi_{0,B})^T$ of $-\mathbf{C}_2$ are given by

$$\phi_{0,A} = \sqrt{\frac{R_{2,B} + K_{BA} - \frac{1}{T_{2,0}^{app}}}{\frac{V_A}{T_{2,1}^{app}} - \frac{V_A}{T_{2,0}^{app}}}}, \tag{50}$$

$$\phi_{0,\mathrm{B}} = \sqrt{\frac{R_{1,\mathrm{A}} + K_{\mathrm{AB}} - \frac{1}{T_{2,0}^{\mathrm{app}}}}{\frac{V_{\mathrm{B}}}{T_{2,1}^{\mathrm{app}}} - \frac{V_{\mathrm{B}}}{T_{2,0}^{\mathrm{app}}}}}, \tag{51}$$

with $V_{\mathrm{A}}|\phi_{0,\mathrm{A}}|^2 + V_{\mathrm{B}}|\phi_{0,\mathrm{B}}|^2 = 1$, and the corresponding eigenvalue is

$$\frac{1}{T_{2,0}^{\mathrm{app}}} = \frac{1}{2}(R_{2,\mathrm{A}} + K_{\mathrm{AB}} + R_{2,\mathrm{B}} + K_{\mathrm{BA}}) - \frac{1}{2}\sqrt{(R_{2,\mathrm{A}} + K_{\mathrm{AB}} - R_{2,\mathrm{B}} - K_{\mathrm{BA}})^2 + 4K_{\mathrm{AB}}K_{\mathrm{BA}}}. \tag{52}$$

The components of the eigenvector $(\phi_{1,\mathrm{A}}, \phi_{1,\mathrm{B}})^{\mathrm{T}}$ of $-\mathbf{C}_2$ are given by

$$\phi_{1,\mathrm{A}} = -\sqrt{\frac{R_{2,\mathrm{A}} + K_{\mathrm{AB}} - \frac{1}{T_{2,0}^{\mathrm{app}}}}{\frac{V_{\mathrm{A}}}{T_{2,1}^{\mathrm{app}}} - \frac{V_{\mathrm{A}}}{T_{2,0}^{\mathrm{app}}}}}, \tag{53}$$

$$\phi_{1,\mathrm{B}} = \sqrt{\frac{R_{2,\mathrm{B}} + K_{\mathrm{BA}} - \frac{1}{T_{2,0}^{\mathrm{app}}}}{\frac{V_{\mathrm{B}}}{T_{2,1}^{\mathrm{app}}} - \frac{V_{\mathrm{B}}}{T_{2,0}^{\mathrm{app}}}}}, \tag{54}$$

with $V_{\mathrm{A}}|\phi_{1,\mathrm{A}}|^2 + V_{\mathrm{B}}|\phi_{1,\mathrm{B}}|^2 = 1$, and the corresponding eigenvalue is

$$\frac{1}{T_{2,1}^{\mathrm{app}}} = \frac{1}{2}(R_{2,\mathrm{A}} + K_{\mathrm{AB}} + R_{2,\mathrm{B}} + K_{\mathrm{BA}}) + \frac{1}{2}\sqrt{(R_{2,\mathrm{A}} + K_{\mathrm{AB}} - R_{2,\mathrm{B}} - K_{\mathrm{BA}})^2 + 4K_{\mathrm{AB}}K_{\mathrm{BA}}}. \tag{55}$$

These eigenvectors also satisfy $\phi_{0,\mathrm{A}}^{*} V_{\mathrm{A}} \phi_{1,\mathrm{A}} + \phi_{0,\mathrm{B}}^{*} V_{\mathrm{B}} \phi_{1,\mathrm{B}} = 0$. The system of ODEs in Eq. (41) can also be solved analytically to give

$$\begin{pmatrix} m_{2,\mathrm{A}}(t_2) \\ m_{2,\mathrm{B}}(t_2) \end{pmatrix} = \exp\{\mathbf{C}_2 t_2\} \begin{pmatrix} m_{2,\mathrm{A}}(0) \\ m_{2,\mathrm{B}}(0) \end{pmatrix} = \sum_{i=0}^{1} \begin{pmatrix} \phi_{i,\mathrm{A}} \\ \phi_{i,\mathrm{B}} \end{pmatrix} \exp\left\{-\frac{t_2}{T_{2,i}^{\mathrm{app}}}\right\} b_i, \tag{56}$$

where $b_i = \phi_{i,\mathrm{A}}^{*} V_{\mathrm{A}} m_{1,\mathrm{A}}(t_1) + \phi_{i,\mathrm{B}}^{*} V_{\mathrm{B}} m_{1,\mathrm{B}}(t_1)$ since $m_{2,\mathrm{A}}(0) = m_{1,\mathrm{A}}(t_1)$ and $m_{2,\mathrm{B}}(0) = m_{1,\mathrm{B}}(t_1)$. The resulting signal is $h(t_1, t_2) \propto V_{\mathrm{A}} m_{2,\mathrm{A}}(t_2) + V_{\mathrm{B}} m_{2,\mathrm{B}}(t_2)$. The analytic amplitude for each peak in the T_1-T_2 correlation with position $\{T_{1,j}^{\mathrm{app}}, T_{2,i}^{\mathrm{app}}\}$ can again be calculated as $A_{ij} = d_i Q_{ij} a_j$ for $i, j = 0$ or 1. Here, $d_i = V_{\mathrm{A}}\phi_{i,\mathrm{A}} + V_{\mathrm{B}}\phi_{i,\mathrm{B}}$ and $Q_{ij} = \phi_{i,\mathrm{A}}^{*} V_{\mathrm{A}} \psi_{j,\mathrm{A}} + \phi_{i,\mathrm{B}}^{*} V_{\mathrm{B}} \psi_{j,\mathrm{B}}$.

For this investigation, relaxation was considered at the surface of pore A, in the fast diffusion limit. No surface relaxation took place in pore B, so $S_{\mathrm{B}}^{\mathrm{relax}} = 0$, though bulk relaxation did occur in both pores as usual. Weakly coupled pores governed by fast diffusion ($\mathcal{D} = 1 \times 10^{-9}\ \mathrm{m^2\,s^{-1}}$) were simulated in 2D as two square pores of dimension $L_{\mathrm{p}} = 10\ \mu\mathrm{m} \times 10\ \mu\mathrm{m}$ connected by a throat of length $L_{\mathrm{t}} = 1\ \mu\mathrm{m}$ and width 0.1 μm [see Fig. 4(b)]. Surface relaxation occurred only on the upper and lower walls of pore A with $\rho_1 = 10\ \mu\mathrm{m\,s^{-1}}$ and $\rho_2 = 50\ \mu\mathrm{m\,s^{-1}}$. The other walls were simply reflective. A second coupled pore simulation (not illustrated) was governed by slow diffusion with the pore dimensions set to 100 $\mu\mathrm{m} \times$ 100 μm and $\rho_1 = 100\ \mu\mathrm{m\,s^{-1}}$, $\rho_2 = 500\ \mu\mathrm{m\,s^{-1}}$. The connecting throat dimensions remained the same. No analytic solution was determined for coupled pores governed by slow diffusion. The simulations generated time-domain signal data to be used as input for inversion to estimate the T_1-T_2 correlation distribution. No noise was added to these simulated data in order to identify any inherent inaccuracies or discrepancies between the simulation and analytic predictions. Inversion of these simulated data was performed separately by non-negative Tikhonov regularization, half-bound Tikhonov regularization, and the primal-dual method.

IV. RESULTS AND DISCUSSION FOR SIMULATED SYSTEMS

We aim to assign peaks in T_1-T_2 correlations for the reconstruction results that follow in this section. The reconstructions were performed from the time-domain signals simulated using the finite-element method. Each peak present in the plots is assigned with a pair of indices, $\{ij\}$, where the i label increases from right to left and the j label increases from top to bottom in each plot. The apparent transverse and longitudinal relaxation times that specify the coordinates of the peak are interpreted to correspond to inverse eigenvalues $\{T_{2,i}^{\mathrm{app}}, T_{1,j}^{\mathrm{app}}\}$. A peak label $\{ij\}$ thus refers to a particular eigenvalue combination. The eigenvalues $\{1/T_{2,i}^{\mathrm{app}}, 1/T_{1,j}^{\mathrm{app}}\}$ are those of either relaxation-diffusion operators or relaxation-coupling matrices. Each peak has an associated amplitude given by A_{ij}. These amplitudes are compared to the available analytic predictions, which were calculated using the known eigenfunctions or eigenvectors. The analytic solutions strictly return delta functions and the peaks shown in such cases are artificially broadened for convenient display.

T_1-T_2 correlations were calculated using Eq. (7) to describe the diffusion of magnetization in an isolated pore for a range of surface relaxivities [see Figs. 5(a)–5(c)]. When $\rho r_{\mathrm{sph}}/\mathcal{D} < 1$ (i.e., fast diffusion), the peak corresponding to the lowest longitudinal and transverse eigenvalues ($i, j = 0$) dominates the distribution [see Fig. 5(a)]. This single peak is present with amplitude $A_{00} = 1$, located at eigenvalues $1/T_{1,0}^{\mathrm{app}} = 0.53\ \mathrm{s^{-1}}$ and $1/T_{2,0}^{\mathrm{app}} = 0.64\ \mathrm{s^{-1}}$, each being a weighted sum of the bulk and surface relaxation rates. In this case, all higher-index peaks ($i, j \geqslant 1$) are absent, having zero amplitude. Moving to the intermediate and slow diffusion regimes [Figs. 5(b) and 5(c)] additional peaks become visible. A total of nine peaks are visible in Fig. 5(c), corresponding to all eigenvalues with indices $i, j \leqslant 2$. Note that the {01}, {02}, and {12} peaks are located at apparent relaxation times (inverse eigenvalues) where $T_{2,i}^{\mathrm{app}} > T_{1,j}^{\mathrm{app}}$ (i.e., $i < j$) and so have negative amplitude, given $T_1^{\mathrm{b}} = T_2^{\mathrm{b}}$. The higher-index peaks ($i, j \geqslant 3$) have too small amplitude to be visible with the scale of the contours. The expected observable peaks are thus those for eigenvalue combinations with $i, j \leqslant 2$.

Time-domain data corresponding to the three diffusion regimes, generated by finite-element simulation, were inverted using non-negative Tikhonov regularization [Figs. 5(d)–5(f)], half-bound Tikhonov regularization

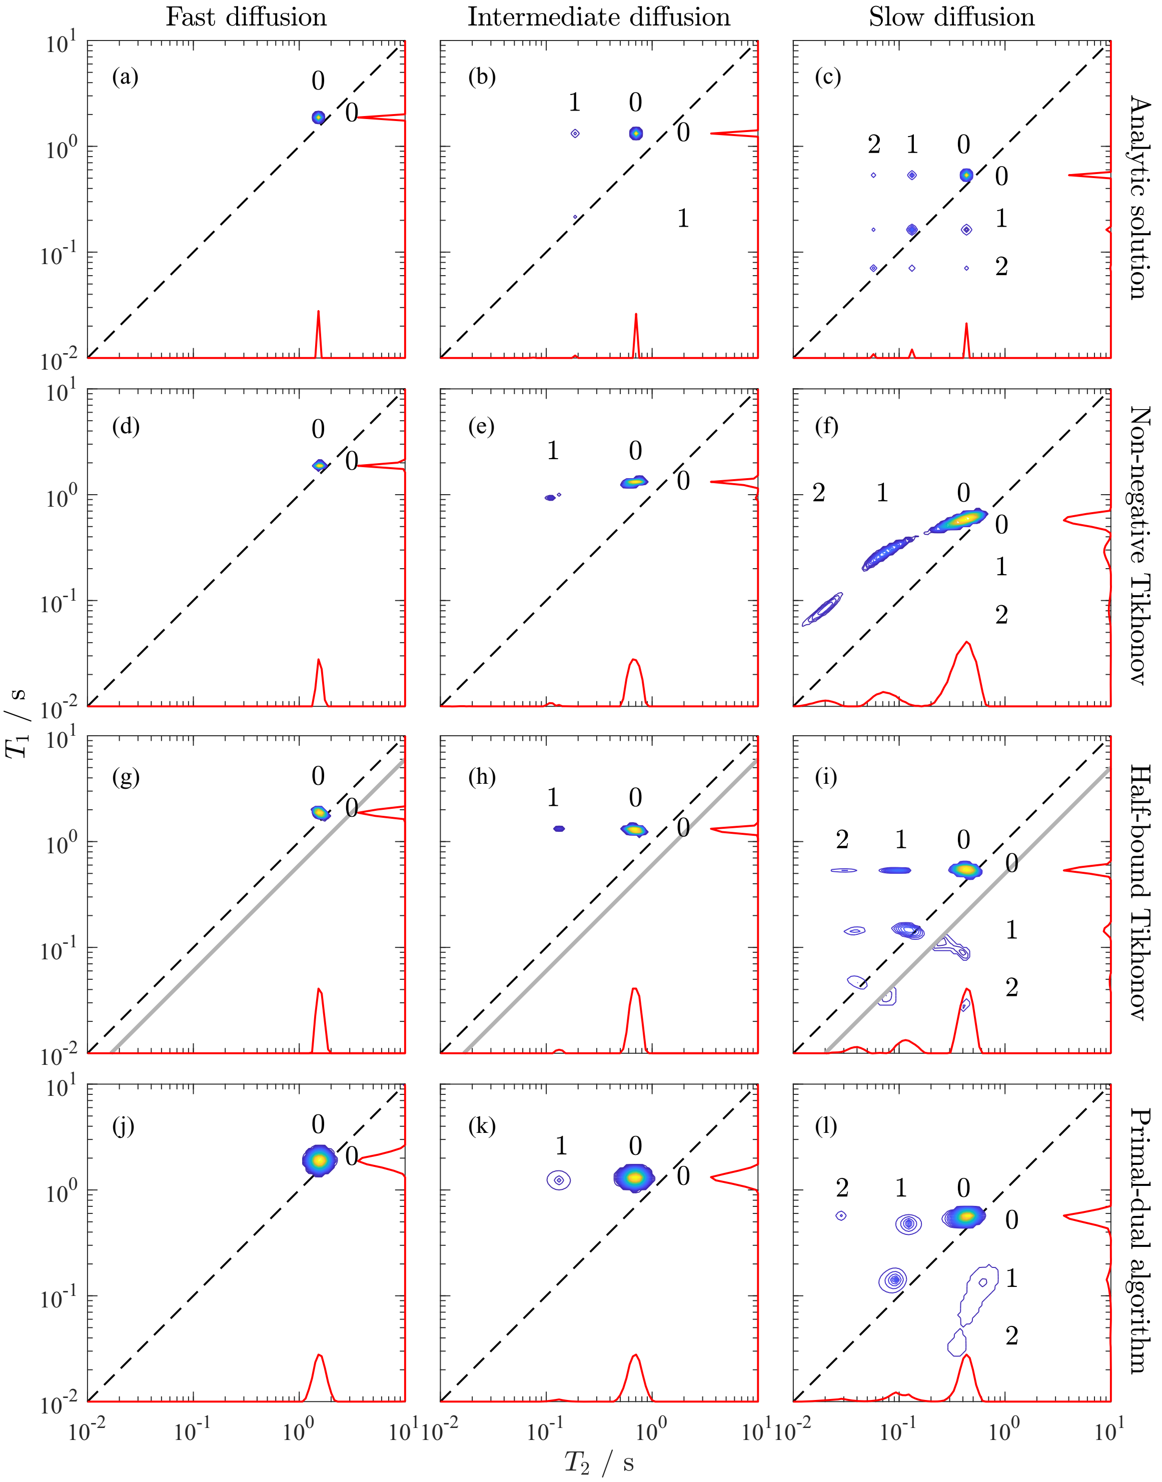

FIG. 5. T_1-T_2 correlation plots for simulated diffusion in an isolated spherical pore illustrated in Fig. 4(a). The Brownstein-Tarr diffusion regime was varied by adjusting the surface relaxivity to give [(a), (d), (g), (j)] fast diffusion, [(b), (e), (h), (k)] intermediate diffusion, and [(c), (f), (i), (l)] slow diffusion. Analytic solutions (a)–(c) were calculated using Eq. (7); broadened distributions are shown for clarity. Finite-element simulations produced time-domain data inverted using (d)–(f) non-negative Tikhonov regularization, (g)–(i) half-bound Tikhonov regularization, and (j)–(l) the primal-dual method. The contour intervals vary between plots, but each 2D plot has been normalized to unit signal (total peak integral). Marginal projections are included for clarity. In each plot, the dashed diagonal line indicates $T_1 = T_2$. Peaks located at $T_2 > T_1$ (i.e., below this diagonal) have negative amplitude. The eigenvalue indices $\{ij\}$ are indicated in each plot with i increasing from right to left and j increasing from top to bottom. For the half-bound Tikhonov regularization (g)–(i), the gray diagonal line indicates the change of sign in the kernel; artifacts have been manually removed (see Fig. 6 for details).

[Figs. 5(g)–5(i)], and the primal-dual method [Figs. 5(j)–5(l)]. For the fast diffusion regime, the 2D distribution (single peak) was entirely positive, so all three inversion methods delivered consistent solutions. The primal-dual method provided a broader peak than Tikhonov regularization due to the weight of the second-derivative penalty (see Sec. II C), chosen for stability in the solution. Reduction in the value of λ_L relative to λ_I may provide narrower peaks at the risk of introducing spurious artifacts.

Moving to the intermediate diffusion regime [Fig. 5(b)], higher-index peaks $i, j, \geqslant 1$ are seen with non-negligible amplitude, although the {00} peak remains dominant. Peak integrals were calculated for all $i, j = 0, 1$, yielding amplitudes $A_{00} = 0.95$, $A_{10} = 0.04$, $A_{11} = 0.01$, $A_{01} = -0.01$; the latter peak is not visible in Fig. 5(b) due to the scale of the contours. Higher-index peaks $i, j \geqslant 2$ had amplitudes of $|A_{ij}| \ll 0.001$ and so were disregarded. The {00} and {10} peaks were recovered from the simulated time-domain data by the half-bound Tikhonov regularization and the primal-dual methods [Figs. 5(h) and 5(k)], with $A_{00} \approx 0.97$ and $A_{10} \approx 0.03$ in both cases. The recovered {10} peaks are slightly shifted from the analytic prediction, but consistent in both inverted solutions, so this variation is attributed to details of the finite-element model. The low-intensity peaks {01} (negative amplitude) and {11} were not recovered by the inversion algorithms. It is interesting to note that the solution achieved by non-negative Tikhonov regularization [Fig. 5(e)] is subtly distorted (the {10} peak is displaced on the T_1 axis), even though the {01} peak has negligible amplitude here. This observation highlights the sensitivity of the non-negative inversion to details of the time-domain data that result from even small deviations from the non-negativity assumption.

When the magnetization evolution is governed by the slow diffusion regime, numerous higher-index peaks ($i, j \leqslant 2$) have significant amplitude [see Fig. 5(c)]. The majority of the signal amplitude is still associated with the {00} peak, such that $A_{00} \sim 0.7$. However, the {01}, {02}, and {12} peaks (negative amplitudes) give $\sum_{i<j\leqslant 2} A_{ij} \approx -0.08$, i.e., are equivalent to 8% of the total signal, and are all visible in Fig. 5(c). Higher-index combinations $i, j \geqslant 3$ with amplitudes $|A_{ij}| \leqslant 0.001$ were calculated but omitted from the plot. The amplitudes of the peaks with $i, j \leqslant 2$ are given in Table I.

The solution obtained by non-negative Tikhonov regularization of the simulated time-domain data is a poor reconstruction of the analytic solution [see Fig. 5(f)]. Three peaks are present in this T_1-T_2 correlation plot, albeit elongated approximately parallel to the $T_1 = T_2$ diagonal. These peaks could be associated with the positions of largest expected amplitude: {00}, {10}, and {11} (see Table I). Assuming this assignment is valid, both the amplitude and position of the {10} and {11} peaks are substantially wrong. The {00} peak has the correct amplitude and is centered approximately on the expected apparent relaxation times (inverse eigenvalues). Alternatively, the three recovered peaks could be assigned to the diagonal index combinations: {00}, {11}, and {22}, in which case the amplitudes are closer to the expected values, but the apparent relaxation times remain incongruent with the analytic solution. It is clear that the presence of peaks with significant negative amplitude distort the solution obtained by non-negative Tikhonov regularization, preventing robust interpretation in terms of apparent relaxation times and signal amplitudes.

TABLE I. Signal (peak) amplitudes extracted from T_1-T_2 correlation plots for relaxation in an isolated spherical pore governed by slow diffusion. The apparent relaxation times correspond to the eigenvalues (rates) from the analytic solution, with the matching indices given in curled braces (i across columns, j over rows). Peaks with indices $i, j = 0 \ldots 8$ were used in the analytic calculation, although only the low-index ($i, j, \leqslant 2$) amplitudes are included here. These data correspond to the plots in Figs. 5(c), 5(f), 5(i), and 5(l) for the analytic, non-negative Tikhonov regularization, half-bound Tikhonov regularization, and primal-dual methods, respectively.

			$T_{2,i}^{\text{app}}$ / ms		
A_{ij}			431 {0}	128 {1}	59 {2}
			Analytic solution		
	549	{0}	0.743	0.059	0.022
	160	{1}	− 0.052	0.125	0.017
	70	{2}	− 0.009	− 0.016	0.037
			Non-negative Tikhonov		
$T_{1,j}^{\text{app}}$ / ms	549	{0}	0.753	0.172	0
	160	{1}	0	0.060	0
	70	{2}	0	0	0
			Half-bound Tikhonov		
	549	{0}	0.785	0.067	0.016
	160	{1}	− 0.049	0.142	0.027
	70	{2}	− 0.004	− 0.017	0.038
			Primal-dual algorithm		
	549	{0}	0.754	0.091	0.025
	160	{1}	− 0.032	0.118	0.020
	70	{2}	− 0.010	0	0.018

Half-bound Tikhonov regularization [Fig. 5(i)] recovers all nine of the expected peaks. The apparent relaxation times (inverse eigenvalues) correspond approximately to the analytic solution for $i > j$, with some distortion along the T_2 axis (similar to the inverted data for the intermediate diffusion regime). The apparent relaxation times for $i < j$ are not rigorously consistent with the analytic solution, but the peaks are sufficiently well separated that robust assignment of i, j is achieved everywhere. The peak amplitudes (see Table I) are reasonable estimates of those in the analytic solution. We note that the half-bound Tikhonov regularization provides the only solution containing all the expected observable peaks. To achieve this solution, a change in sign was enforced at $T_2/T_1 = 0.4$ and artifacts generated along the diagonal were manually removed. This manual editing stage is illustrated in Fig. 6. The artifacts are more significant when genuine peaks are present near the discontinuity in sign. The half-bound Tikhonov regularization was repeated with the change in sign located at different ratios of T_2/T_1 in the range 0.8 to 0.3. The result with the ratio being 0.4, shown in Fig. 6(a), contained the minimum artifactual signal amplitude, even though it appears the peaks at {12} and {11} have been truncated by the discontinuity. Notwithstanding this apparent limitation, it is worth noting that the artifacts in Fig. 6(a) do not observably contribute to the marginal projections and the 1D T_1 and

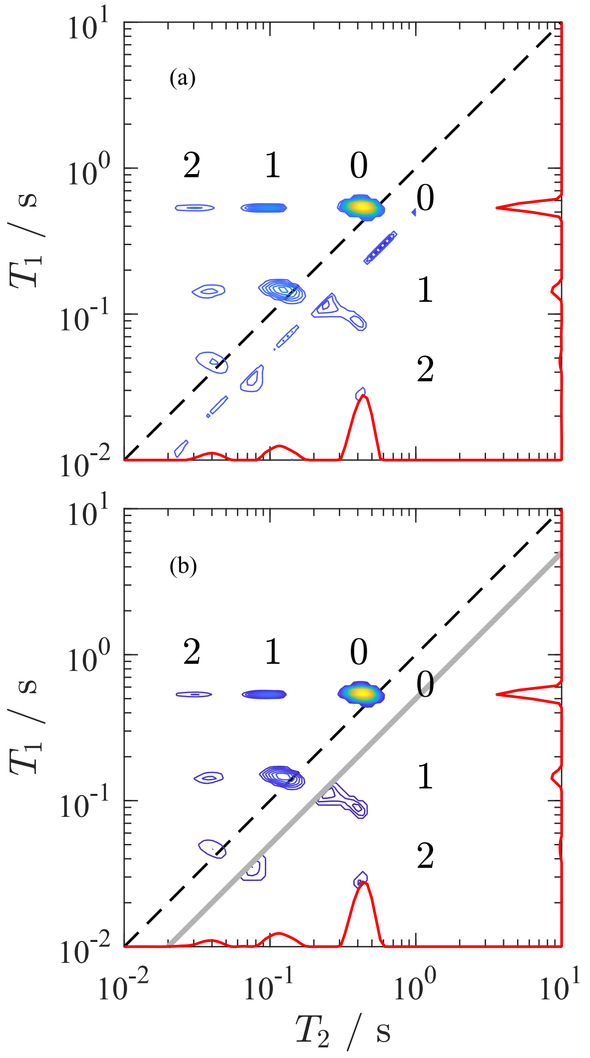

FIG. 6. T_1-T_2 correlation plots for simulated diffusion in an isolated spherical pore (slow diffusion regime) reconstructed using half-bound Tikhonov regularization. The original solution (a) contains artifactual peaks along the $T_2/T_1 = 0.4$ diagonal, associated with the discontinuity in sign imposed on the kernel. Peaks located below this diagonal have negative amplitude. The artifacts have been manually removed in (b), where $T_2/T_1 = 0.4$ is indicated by the gray diagonal line; this plot is identical to Fig. 5(i). The dashed diagonal line indicates $T_1 = T_2$, and the eigenvalue indices $\{ij\}$ are indicated in each plot with i increasing from right to left and j increasing from top to bottom.

T_2 distributions are visually unaltered by the manual editing process [cf. Fig. 6(b)].

Finally, we move to Fig. 5(l) to consider the solution achieved by the primal-dual inversion method. The primal-dual algorithm recovers the expected observable negative-amplitude peaks, with the exception of $\{12\}$, and largely succeeds in returning the correct amplitudes (Table I) and apparent relaxation times for all the other combinations. Peaks not visible in Fig. 5(l), due to the scale of the contours, are present in the solution matrix (**X**). Some distortion of the solution is apparent, e.g., the $\{11\}$, $\{02\}$, and $\{20\}$ peaks are shifted to shorter T_2 times compared to the analytic prediction, but are largely consistent with the half-bound Tikhonov regularization solution. We conjecture that the $\{12\}$ peak has been omitted due the thresholding implicit in the primal-dual method, and suggest this is a limitation of the method. Nevertheless, the solution is considered a good approximation to the analytic distribution.

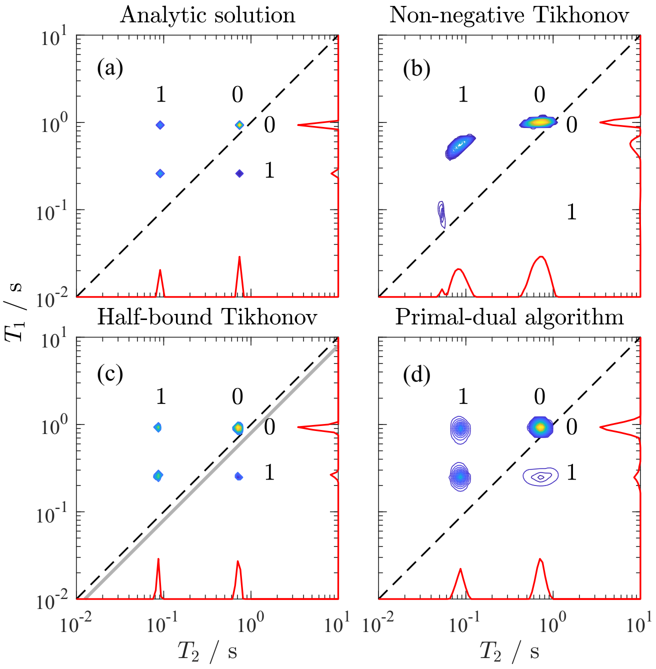

FIG. 7. T_1-T_2 correlation plots for simulated diffusion in two weakly coupled pores illustrated in Fig. 4(b). The choice of ρ, L, and $\mathcal{D}$ corresponded to fast diffusion in each pore. An analytic solution (a) was calculated using Eqs. (39)–(42). Finite-element simulations produced time-domain data which were inverted using (b) non-negative Tikhonov regularization, (c) half-bound Tikhonov regularization, and (d) the primal-dual method. The contour intervals vary between plots, but each 2D plot has been normalized to unit volume (total peak integral). Marginal projections are included for clarity. The eigenvalue indices $\{ij\}$ are indicated in each plot with i increasing from right to left and j increasing from top to bottom. The dashed diagonal line indicates $T_1 = T_2$, and the gray diagonal line at $T_2/T_1 = 0.8$ in (c) indicates the change in sign imposed on the kernel. Peaks located at $T_2 > T_1$ (i.e., $i < j$) have negative amplitude.

The analytic T_1-T_2 correlation for two weakly coupled pores in fast diffusion is shown in Fig. 7(a). Only four eigenfunction combinations exist ($i, j = 0, 1$), and the majority of the signal amplitude resides in the diagonal $\{00\}$ and $\{11\}$ peaks (see Table II). The $\{01\}$ peak has negative amplitude, and accounts for approximately 9% of the total signal. Note that the marginal projections are positive always [this was also true for the isolated pore results in Figs. 5(a)–5(c), but less obvious due to the larger ($i, j = 0 \ldots \infty$) range of valid eigenfunction combinations]. Again, time-domain data were generated by finite-element simulation and inverted using the three reconstruction methods. The solution recovered by non-negative Tikhonov regularization [Fig. 7(b)] is a very poor reconstruction of the analytic distribution. The amplitude and apparent relaxation times of the $\{00\}$ peak are reasonable,

TABLE II. Signal (peak) amplitudes extracted from T_1-T_2 correlation plots for two weakly coupled pores governed by fast diffusion. The relaxation times correspond to the eigenvalues (rates) from the analytic solution, with the matching indices given in curled braces (i across columns, j over rows). These data correspond to the plots in Figs. 7(a)–7(d) for the analytic, non-negative Tikhonov regularization, half-bound Tikhonov regularization, and primal-dual methods, respectively.

			$T_{2,i}^{\mathrm{app}}$ / ms	
	A_{ij}		714 {0}	86 {1}
			Analytic solution	
	921	{0}	0.684	0.170
	256	{1}	– 0.086	0.232
			Non-negative Tikhonov	
$T_{1,j}^{\mathrm{app}}$ / ms	921	{0}	0.628	0.346
	256	{1}	0	0.026
			Half-bound Tikhonov	
	921	{0}	0.690	0.163
	256	{1}	– 0.078	0.224
			Primal-dual algorithm	
	921	{0}	0.686	0.174
	256	{1}	– 0.083	0.225

but the amplitudes are wrong with $A_{10} > A_{11}$. The apparent longitudinal relaxation times are too low for both {10} and {11} peaks; the apparent transverse relaxation time is also wrong for the {1, 1} peak in particular. The {01} peak is absent as expected due to the non-negativity constraint. Assignment of the peaks is robust insomuch as there is only one reasonable choice. However, interpretation of the amplitudes of the peaks remains ambiguous due to the substantial distortion of the solution. The half-bound Tikhonov regularization and primal-dual methods are able to recover excellent representations of the analytic distribution, Figs. 7(c) and 7(d), respectively. For the half-bound regularization, the change in sign was imposed at $T_2/T_1 = 0.8$ and produced negligible artifacts in this case. The peaks recovered by the primal-dual algorithm are substantially broadened, yet the eigenvalues and amplitudes (Table II) match the analytic result.

To demonstrate the application of the half-bound Tikhonov and primal-dual algorithms to more realistic systems, the evolution of magnetization in two weakly coupled pores governed by slow diffusion was simulated. The T_1-T_2 correlation obtained by non-negative Tikhonov regularization of the time-domain data is shown in Fig. 8(a). Four peaks are identified. The peak at longest relaxation times $T_1 = T_2 = 2$ s corresponds to bulk relaxation in pore B (no surface relaxation). The three peaks at shorter relaxation times are attributed primarily to slow diffusion in pore A with some mixing between the two pores. Any peaks with apparent $T_2 > T_1$ are suppressed by the non-negativity constraint. The solution achieved using half-bound Tikhonov regularization [Fig. 8(b)] reveals additional peaks attributed to slow diffusion eigenvalues. As well as the peak at $T_1 = T_2 = 2$ s, six other components with positive amplitude are identified in the plot, consistent with the slow diffusion isolated pore result in Fig. 5(c). In Fig. 8(b), two components with negative ampli-

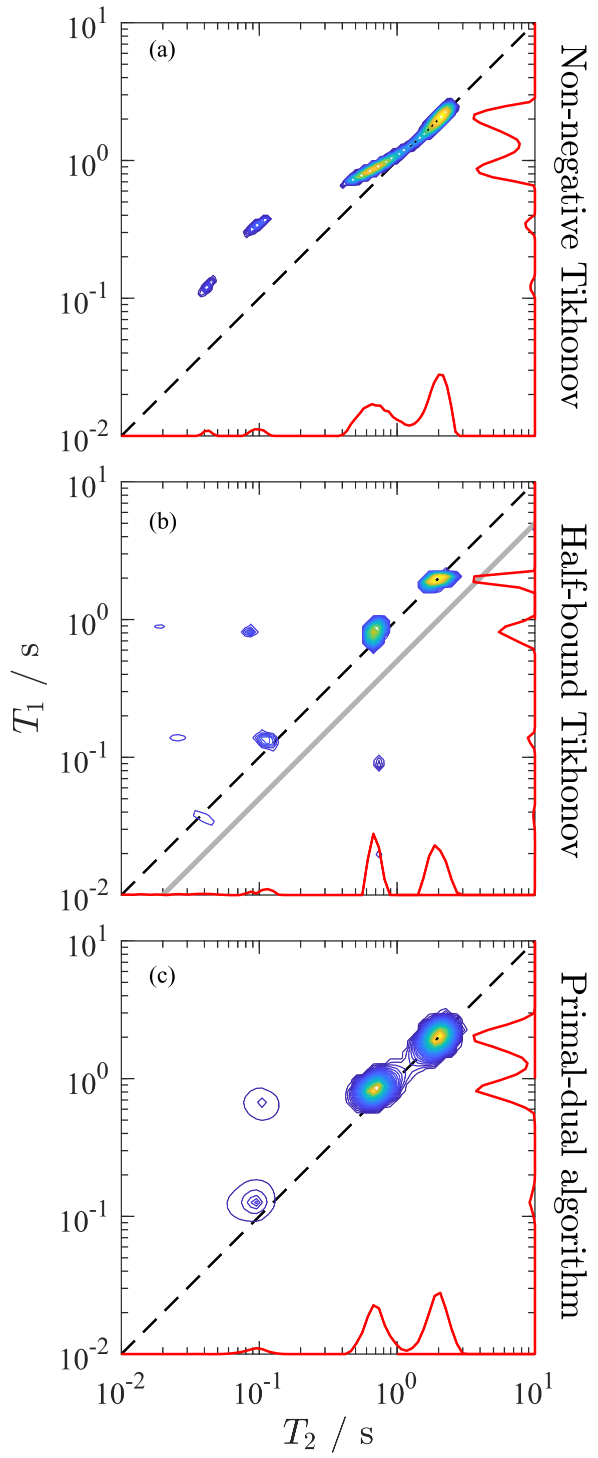

FIG. 8. T_1-T_2 correlation plots for simulated diffusion in two weakly coupled pores. The choice of ρ, L, and $\mathcal{D}$ corresponded to slow diffusion in each pore. Finite-element simulations produced time-domain data inverted using (a) non-negative Tikhonov regularization, (b) half-bound Tikhonov regularization, and (c) the primal-dual method. The contour intervals vary between plots, but each 2D plot has been normalized to unit volume (total peak integral). Marginal projections are included for clarity. In each plot, the dashed diagonal line indicates $T_1 = T_2$. Peaks located at $T_2 > T_1$ (i.e., below the dashed diagonal) have negative amplitude. The gray diagonal line at $T_2/T_1 = 0.5$ in (b) indicates the change in sign imposed on the kernel.

tude are visible, both with $T_2 = 0.7$ s (the low-amplitude component at $T_1 = 0.02$ s is obscured by the marginal distribution plot). A negative-amplitude component was thus expected at $T_2 = 0.1$ s, $T_1 = 0.02$ s but may have been lost amongst the artifacts along the $T_2/T_1 = 0.5$ diagonal that were removed. The distribution obtained using the primal-dual method is shown in Fig. 8(c). Four peaks are identified, essentially consistent with the non-negative Tikhonov regularization solution in Fig. 8(a). Although the peaks are broadened by our choice of primal-dual algorithm, they remain distinct in the marginal projections and the apparent relaxation times approximately match those obtained by the half-bound Tikhonov regularization. The Tikhonov regularization algorithm tends to preserve low-intensity features, whereas the primal-dual algorithm tends to preserve high-intensity features, and this difference in reconstruction bias is evident in these distributions. Despite the loss of low-amplitude features in Fig. 8(c), the principal peaks are present and permit robust interpretation in terms of amplitude and apparent relaxation times. The solutions provided by half-bound Tikhonov regularization and the primal-dual method highlight the complexity that can exist in even relatively simple coupled pore systems.

Under favorable circumstances, the primal-dual algorithm recovers a T_1-T_2 correlation containing peaks of negative amplitude. Applied to simulated data, the results were very encouraging, with peak amplitudes and positions recovered correctly. Low-intensity features were occasionally rejected due to the shrinkage effect of the ℓ_1 penalties, yet the recovered distributions were good approximations to the analytic results where available. However, these simulations were essentially free from random noise or systematic errors (e.g., interference sources) that are present in actual measurements. In the next section, we apply the inversion method to experimental data.

Half-bound Tikhonov regularization provides solutions with narrow peaks, simplifying interpretation of the distribution. Successful application of this method relies on (1) peaks with negative amplitude having very different positions relative to the positive-amplitude contributions, and (2) the range of apparent relaxation times where the peaks have negative amplitude being known *a priori*, e.g., where apparent $T_2 > T_1$. In the simple cases of coupled pores or Brownstein-Tarr slow diffusion presented here, these requirements are met, but may not be satisfied in general. There remains an implementation challenge of defining the region of nonpositive amplitude, plus manual editing of the solution to remove artifacts associated with the imposed discontinuity in sign. We know from experience that methods based on Tikhonov regularization are robust in application to noisy experimental measurements.

V. EXPERIMENTAL DEMONSTRATION

So far we have tested the half-bound Tikhonov regularization and primal-dual algorithms against simulated signals known to contain relaxation-time components of positive and negative amplitude. Here, we present a single application of these methods to experimental measurements.

A microporous silica powder of particle size between 200 and 500 μm and nominal pore size 4 nm was saturated with deionized water. The wet slurry was centrifuged to sediment

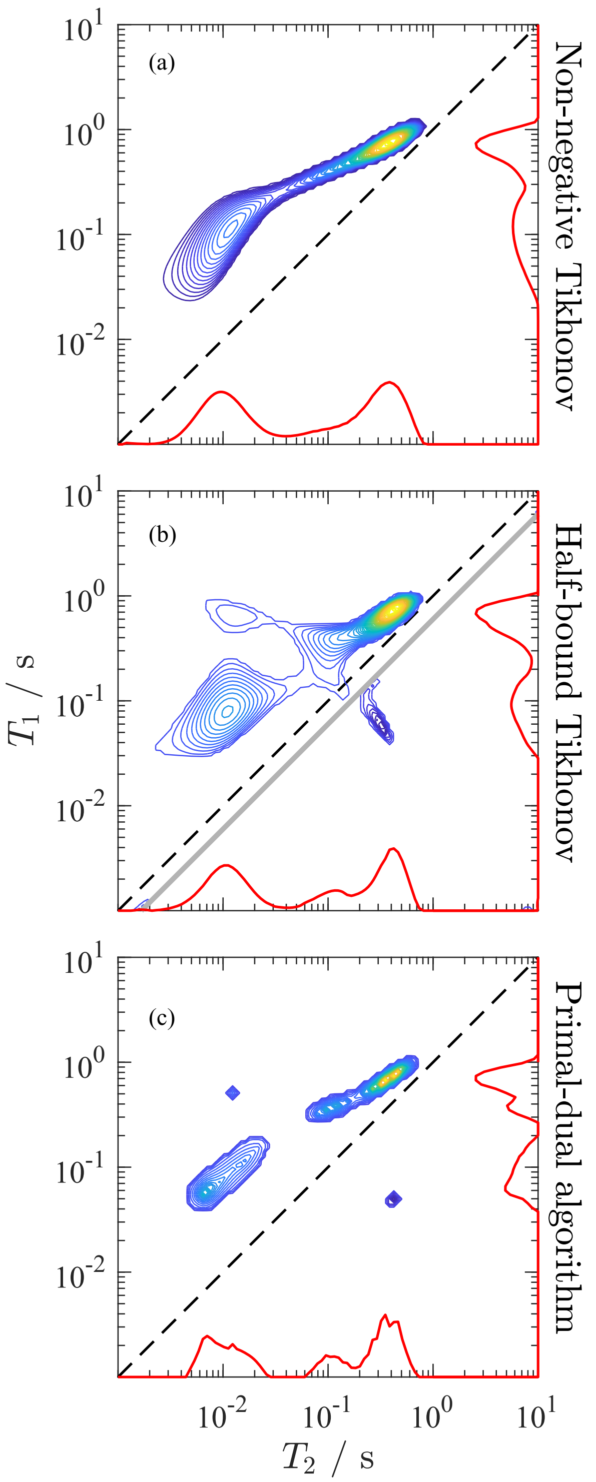

FIG. 9. T_1-T_2 correlation plots for water saturated microporous silica. The time-domain data were inverted using (a) non-negative Tikhonov regularization, (b) half-bound Tikhonov regularization, and (c) the primal-dual method. The contour intervals vary between plots, but each 2D plot has been normalized to unit volume (total peak integral). Marginal projections are included for clarity. In each plot, the dashed diagonal line indicates $T_1 = T_2$. Peaks located at $T_2 > T_1$ (i.e., below the dashed diagonal) have negative amplitude. The gray diagonal line at $T_2/T_1 = 0.6$ in (b) indicates the change in sign imposed on the kernel.

the particles and the excess liquid aspirated from the sample. Diffusive exchange of water was expected between the interparticle and intraparticle porosity. A 2D data set was acquired on an Oxford Instruments bench top magnet ($B_0 = 0.3$ T corresponding to $f_0 = 12.9$ MHz for ^{1}H). Inversion recovery delays were incremented logarithmically from $\tau_1 = 1$ ms to 10 s in 32 steps. CPMG echo trains were acquired with an echo time $t_e = 400$ μs and 6250 echoes. Scans for 8 repetitions were co-added to accommodate the phase cycle for the radio frequency (rf) pulses. The measured data had SNR ≈ 1000.

The T_1-T_2 correlations obtained from the experimental data are shown in Fig. 9. We do not offer any interpretation of the plots here; these results are presented only to demonstrate the applicability of the inversion methods to actual NMR measurements. The non-negative Tikhonov regularization generated a single elongated feature in the T_1-T_2 correlation plot, Fig. 9(a). No distinct off-diagonal peaks are resolvable and the marginal projections are bimodal in both dimensions. A very different plot was obtained using the half-bound Tikhonov regularization, Fig. 9(b). At least two relaxation-time components are distinguishable parallel to the $T_1 = T_2$ diagonal and a pair of asymmetric off-diagonal peaks are present and distinct; the peak below the $T_2/T_1 = 0.6$ diagonal (gray line) has negative amplitude. The T_1 marginal projection appears bi-modal whereas three peaks are visible in the T_2 marginal projection. A consistent result was obtained using the primal-dual method, Fig. 9(c). Low-amplitude features were suppressed by the algorithm, as expected, so that three peaks are visible parallel to the $T_1 = T_2$ diagonal and in both marginal projections. Again, a pair of asymmetric off-diagonal peaks are present and distinct with the peak below the $T_1 = T_2$ diagonal having negative amplitude. It is clear from these plots that both the half-bound Tikhonov regularization and primal-dual method are suitable for use on experimental data containing noise and systematic errors (in this case, imperfect inversion of the magnetization). Both methods reveal the detailed structure of the T_1-T_2 distribution that is lost when using non-negative Tikhonov regularization.

VI. CONCLUSIONS

NMR relaxation-time correlations are powerful tools for studying porous media, with the T_1-T_2 correlation being popular. A non-negative Tikhonov regularization algorithm is often used to invert the first-kind Fredholm integral equation describing the time-domain data, and in most cases this approach is satisfactory. However, diffusive transport in conjunction with differences in longitudinal and transverse surface relaxation can result in the occurrence of peaks with positive and negative amplitude in the 2D distribution. Mathematical models of such situations yield eigenvalues that determine the apparent relaxation times and eigenfunctions (or eigenvectors) that determine the amplitudes of the peaks. In pore systems governed by the Brownstein-Tarr "fast" diffusion regime, peaks associated with high eigenvalues (low apparent relaxation times) have negligible amplitude and can be safely ignored. In other cases, such as pores governed by the "intermediate" or "slow" diffusion regimes, or for diffusive coupling between a hierarchy of pore sizes (or spatially localized surface relaxivities), the peaks with negative amplitude are significant. We have demonstrated that, under such conditions, the usual non-negative Tikhonov regularization approach to inverting the time-domain data produces a distorted solution, and robust interpretation (i.e., quantitative estimation of apparent relaxation times and signal amplitudes) is not feasible. Two numerical inversion methods capable of recovering all the expected observable peaks in a T_1-T_2 correlation experiment were evaluated: (1) half-bound Tikhonov regularization with nonpositivity enforced over a limited range of T_1 and T_2 with non-negativity elsewhere, and (2) a primal-dual method for a generalized LASSO that contains no assumptions regarding the sign of the solution. The half-bound Tikhonov regularization provided qualitatively superior distributions (narrow peaks and accurate apparent relaxation times) for the simulated situations at the expense of artifacts in the solution that required manual elimination. This inversion method is robust in application to noisy experimental data, so long as the region of the T_1-T_2 correlation space corresponding to peaks of negative amplitude is known, and can be recommended for T_1-T_2 correlation experiments. For general inversions, the primal-dual algorithm is recommended and may be used to define the apparent relaxation time ranges for enforcing nonpositivity in the half-bound Tikhonov regularization. The primal-dual method yielded the clearest result for the experimental case considered.

ACKNOWLEDGMENTS

We thank E. Fordham and A. Clarke of Schlumberger Cambridge Research for assistance with the COMSOL simulations.

[1] J. Mitchell, L. Gladden, T. Chandrasekera, and E. Fordham, Prog. Nucl. Magn. Reson. Spectrosc. **76**, 1 (2014).
[2] Y. Q. Song, L. Venkataramanan, M. D. Hürlimann, M. Flaum, P. Frulla, and C. Straley, J. Magn. Reson. **154**, 261 (2002).
[3] A. Valori, G. Hursan, and S. Ma, Petrophysics **58**, 352 (2017).
[4] K. Washburn and J. Birdwell, J. Magn. Reson. **233**, 17 (2013).
[5] M. Fleury and M. Romero-Sarmiento, J. Petrol. Sci. Eng. **137**, 55 (2016).
[6] P. J. McDonald, J. P. Korb, J. Mitchell, and L. Monteilhet, Phys. Rev. E **72**, 011409 (2005).
[7] Y. Q. Song, Prog. Nucl. Magn. Reson. Spectrosc. **55**, 324 (2009).
[8] D. W. Weber, J. Mitchell, J. McGregor, and L. F. Gladden, J. Phys. Chem. C **113**, 6610 (2009).
[9] G. Lui, Y. Li, and J. Jonas, J. Chem. Phys. **95**, 6892 (1991).
[10] I. Foley, S. Farooqui, and R. Kleinberg, J. Magn. Reson., Ser. A **123**, 95 (1996).
[11] C. D'Agostino, J. Mitchell, M. Mantle, and L. Gladden, Chem. Eur. J. **20**, 13009 (2014).

[12] H. McConnell, J. Chem. Phys. **28**, 430 (1958).

[13] J. H. Lee, C. Labadie, C. S. Springer, and G. S. Harbison, J. Am. Chem. Soc. **115**, 7761 (1993).

[14] K. E. Washburn and P. T. Callaghan, Phys. Rev. Lett. **97**, 175502 (2006).

[15] T. C. Chandrasekera, J. Mitchell, E. J. Fordham, L. F. Gladden, and M. L. Johns, J. Magn. Reson. **194**, 156 (2008).

[16] L. Monteilhet, J. P. Korb, J. Mitchell, and P. J. McDonald, Phys. Rev. E **74**, 061404 (2006).

[17] M. van Landeghem, A. Haber, J. B. d'Espinose de Lacaillerie, and B. Blümich, Concept. Magn. Reson. **36A**, 153 (2010).

[18] Y.-Q. Song, L. Zielinski, and S. Ryu, Phys. Rev. Lett. **100**, 248002 (2008).

[19] L. M. Schwartz, D. L. Johnson, J. Mitchell, T. C. Chandrasekera, and E. J. Fordham, Phys. Rev. E **88**, 032813 (2013).

[20] H. Torrey, Phys. Rev. **104**, 563 (1956).

[21] R. L. Vold, J. S. Waugh, M. P. Klein, and D. E. Phelps, J. Chem. Phys. **48**, 3831 (1968).

[22] H. Y. Carr and E. M. Purcell, Phys. Rev. **94**, 630 (1954).

[23] S. Meiboom and D. Gill, Rev. Sci. Instrum. **29**, 668 (1958).

[24] K. R. Brownstein and C. E. Tarr, Phys. Rev. A **19**, 2446 (1979).

[25] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-posed Problems* (V. H. Winston & Sons, Washington, D.C., 1977).

[26] L. Venkataramanan, Y. Q. Song, and M. D. Hürlimann, IEEE Trans. Signal Process. **50**, 1017 (2002).

[27] J. P. Butler, J. A. Reeds, and S. V. Dawson, SIAM J. Numer. Anal. **18**, 381 (1981).

[28] J. Mitchell, T. Chandrasekera, and L. Gladden, Prog. Nucl. Magn. Reson. Spectrosc. **62**, 34 (2012).

[29] D. Bytchenkoff and S. Rodts, J. Magn. Reson. **208**, 4 (2011).

[30] D. L. Johnson and L. M. Schwartz, Phys. Rev. E **90**, 032407 (2014).

[31] Y.-Q. Song, G. Carneiro, L. M. Schwartz, and D. L. Johnson, Phys. Rev. Lett. **113**, 235503 (2014).

[32] R. Tibshirani, J. R. Statist. Soc. B **58**, 267 (1996).

[33] R. Tibshirani, J. R. Statist. Soc. B **73**, 273 (2011).

[34] M. Figueiredo, R. Nowak, and S. Wright, IEEE J. Sel. Top. Signal Process. **1**, 586 (2007).

[35] A. Stern, D. Donoho, and J. Hoch, J. Magn. Reson. **188**, 295 (2007).

[36] A. Beck and M. Teboulle, SIAM J. Imaging Sci. **2**, 183 (2009).

[37] X. Zhou, G. Su, L. Wang, S. Nie, and X. Ge, J. Magn. Reson. **275**, 46 (2017).

[38] A. Chambolle and T. Pock, J. Math. Imaging Vision **40**, 120 (2011).

[39] A. Reci, A. Sederman, and L. Gladden, J. Magn. Reson. **281**, 188 (2017).

[40] A. Reci, A. Sederman, and L. Gladden, J. Magn. Reson. **284**, 39 (2017).

[41] H. Wang, M. Zhao, J. L. Ackerman, and Y. Song, J. Magn. Reson. **274**, 137 (2017).

[42] J. D. Griffith, J. Mitchell, A. E. Bayly, and M. L. Johns, J. Phys. Chem. B **113**, 7156 (2009).

[43] J. Mitchell and E. Fordham, Microporous Mesoporous Mater. **269**, 109 (2017).

[44] P. C. Hansen, *Discrete Inverse Problems: Insight and Algorithms* (SIAM, Philadelphia, 2010).

[45] E. J. Fordham, L. Venkataramanan, J. Mitchell, and A. Valori, Diffusion Fundamentals (online) **29**, Paper 2, 1 (2017).

[46] S. Boyd and L. Vandenberghe, *Convex Optimization, With Corrections 2008* (Cambridge University Press, Cambridge, 2004).

[47] E. Chouzenoux, S. Moussaoui, J. Idier, and F. Mariette, IEEE Trans. Signal Process. **58**, 6040 (2010).

[48] G. Wahba, SIAM J. Numer. Anal. **14**, 651 (1977).

[49] J. D. Wilson, J. Mater. Sci. **27**, 3911 (1992).

[50] R. J. Tibshirani and J. Taylor, Ann. Statist. **39**, 1335 (2011).

[51] K. Bredies, K. Kunisch, and T. Pock, SIAM J. Imaging Sci. **3**, 492 (2010).

[52] F. Knoll, K. Bredies, T. Pock, and R. Stollberger, Magn. Reson. Med. **65**, 480 (2011).

[53] T. Park and G. Casella, J. Am. Stat. Assoc. **103**, 681 (2008).

[54] D. S. Grebenkov, Rev. Mod. Phys. **79**, 1077 (2007).

[55] D. S. Grebenkov, Concept Magn. Reson. **32A**, 277 (2008).