

BAYESIAN INFERENCE FOR NMR SPECTROSCOPY WITH APPLICATIONS TO CHEMICAL QUANTIFICATION

BY ANDREW GORDON WILSON*, YUTING WU*, DANIEL J. HOLLAND*,
SEBASTIAN NOWOZIN[†], MICK D. MANTLE*, LYNN F. GLADDEN*, ANDREW BLAKE[†]

Cambridge University and Microsoft Research[†]*

Nuclear magnetic resonance (NMR) spectroscopy exploits the magnetic properties of atomic nuclei to discover the structure, reaction state and chemical environment of molecules. We propose a probabilistic generative model and inference procedures for NMR spectroscopy. Specifically, we use a weighted sum of trigonometric functions undergoing exponential decay to model free induction decay (FID) signals. We discuss the challenges in estimating the components of this general model – amplitudes, phase shifts, frequencies, decay rates, and noise variances – and offer practical solutions. We compare with conventional Fourier transform spectroscopy for estimating the relative concentrations of chemicals in a mixture, using synthetic and experimentally acquired FID signals. We find the proposed model is particularly robust to low signal to noise ratios (SNR), and overlapping peaks in the Fourier transform of the FID, enabling accurate predictions (e.g., 1% sensitivity at low SNR) which are not possible with conventional spectroscopy (5% sensitivity).

1. Introduction. Nuclear magnetic resonance (NMR) spectroscopy has greatly advanced our understanding of molecular properties, and is now widespread in analytical chemistry. The theory of nuclear magnetic resonance postulates that protons and neutrons behave like gyroscopes that spin about their axes, generating their own small magnetic fields. These concepts were first described by Rabi et al. (1939), for which Isidor Rabi was awarded the 1944 Nobel prize in physics. Later, Bloch, Hanson and Packard (1946) and Purcell, Torrey and Pound (1946) showed how NMR could be used to understand the structure of molecules in liquids and solids, for which they shared the Nobel prize in physics in 1952. Richard Ernst then won the 1991 Nobel prize in chemistry for developing Fourier transform NMR spectroscopy¹, which led to the prevalence of NMR as an analytic technique. NMR spectroscopy is well suited to studying both organic and inorganic molecules, including proteins, and other biochemical species (Barrett et al., 2013), and is routinely used to identify the structure of unknown chemical species or the composition of mixtures.

NMR spectroscopy is quantitative, chemically specific and non-invasive and therefore can

MSC 2010 subject classifications: 92E99, 60-04, 60-08

Keywords and phrases: nuclear magnetic resonance, spectroscopy, chemical quantification, free induction decay

¹See Ernst (1992) for a comprehensive review.

be used to study molecules *in situ* (Gladden, 1994). A major limitation of conventional NMR spectroscopy is its low sensitivity, which has led researchers to develop ever more powerful magnets to amplify the signal (Barrett et al., 2013). However, stronger magnetic fields have limited returns: they are costly and also impractical in many applications, e.g., plant measurements in the chemical industry (Dalitz et al., 2012). There are therefore strong drivers to develop alternative techniques that can provide chemical information from relatively poor quality data.

In this paper, we propose an alternate approach to NMR spectroscopy. In our approach, we model the free induction decay (FID) signal as a realisation from a probabilistic generative model, and we use Bayesian inference and likelihood principles to infer latent variables in this model. This model allows us to leverage additional information in the FID over conventional Fourier transform spectroscopy, with explicit models of signal decay, phase shifts, frequencies, and noise, and prior information, e.g. about resonant frequencies.

In a body of pioneering work (Bretthorst, 1990; Evilia, Effiong and Whittenburg, 1993; Dou and Hodgson, 1995; Andrec and Prestegard, 1998; Rubtsov and Griffin, 2007; Hutton et al., 2009; Aboutanios, Kopsinis and Rubtsov, 2012), discussed further in section 6, various statistical models have been proposed to model the FID directly. These models, while promising, have not yet been widely adopted, with practitioners typically preferring conventional Fourier transform spectroscopy. We further pursue the line of research instigated by Bretthorst (1990), to develop a statistical model of the FID which we show can be used as a highly practical tool. In particular, we introduce new modelling parameters (such as time delays), new estimation procedures, detailed stress tests, detailed comparisons with conventional methods, implementational details necessary for good performance, and the application to chemical quantification. In general, our approach is robust to noise, and allows one to study systems in which there is significant overlap between peaks in the Fourier transform of the FID signal, and a large number of resonant frequencies. While most approaches to NMR spectroscopy are applied in the high resolution ^1H spectrum, we focus on examples in the ^{13}C spectrum, which is 10^4 times less sensitive.

We begin with a brief introduction to conventional Fourier transform spectroscopy in section 2. We then introduce our alternative model for free induction decay in section 3, and discuss inference and parameter estimation in this model in sections 4 and 5. We resolve general practical difficulties with parameter estimation in statistical models of free induction decay – difficulties which have not been discussed at length in prior work. In particular, estimating a large number of resonant frequencies, jointly with phase parameters and other variables, requires careful navigation of a highly multimodal likelihood surface.

In our experiments of sections 7 and 8, we compare conventional Fourier transform spectroscopy with the proposed method for predicting the relative concentrations of chemicals in a mixture, using both synthetic and experimentally acquired FID signals. We show that the proposed model enables accurate predictions which are not possible with conventional Fourier transform spectroscopy.

2. Conventional Spectroscopy. In this section we outline a conventional Fourier transform spectroscopy approach (Keeler, 2011; Malz and Jancke, 2005). Then, in the following sections, we will propose an alternative method for NMR spectroscopy, and compare both the conventional and proposed approach for quantifying relative chemical concentrations in a mixture.

To understand how NMR spectroscopy works at a high level, imagine a chemical mixture is placed in a strong magnetic field. Nuclei within this mixture that have non-zero spin will interact with the magnetic field to produce “magnetic moments”. We can imagine these magnetic moments as small bar magnets. Radio frequency (rf) pulses are then directed at the mixture, exerting a torque on these “bar magnets”, causing them to precess perpendicular to the strong magnetic field. The rotating magnets create a magnetic flux through a coil in the NMR machine, which induces a current with periodic components at the resonant frequencies of the chemical mixture. The voltage associated with this current is what is measured in an NMR experiment and is called the free induction decay (FID). The resonant frequencies are sensitive to the local molecular structure, which is what permits NMR to be used as a spectroscopic technique.

In conventional Fourier transform spectroscopy, the FID signal is assumed to be generated by two channels, exactly $\pi/2$ out of phase, with each channel perfectly modelled as a noise free sum of cosines at different frequencies $\omega_i^{(j)}$ weighted by intensities $B_i^{(j)}$. Thus taking the Fourier transform of the FID, in principle, would result in a series of spikes (delta functions) at the resonant frequencies, and the relative magnitudes of these spikes (after adjusting for the intensities $B_i^{(j)}$) would be the relative concentrations of chemicals in a mixture. However, in actuality, there is decay in the signal due to, e.g., variations in the magnetic field (Keeler, 2011) (see also section 3 for more detail). Moreover, the trigonometric terms are not in phase, and there is noise in the signal. In section 3, we explicitly model these properties of the FID.

Throughout our experiments we will consider mixtures of cyclohexane and butanone, so we will exemplify the variant of conventional Fourier transform spectroscopy we use in this context; for reference, we follow the general conventional procedure outlined in Keeler (2011) and Malz and Jancke (2005). Table 1 gives the theoretical intensities $B_i^{(j)}$ and frequencies $\omega_i^{(j)}$ for cyclohexane and butanone². Because the NMR spectrometer can only excite spins within a certain resonance frequency range, the extent to which each chemical group is excited differs from each other slightly, and this discrepancy is described by what is known as an *excitation profile*. This phenomenon alters the known intensities $B_i^{(j)}$. We therefore run a calibration experiment to measure the excitation profile, analogous to slice excitation measurements in magnetic resonance imaging (Haacke et al., 1999). We then adjust the theoretical intensity values accordingly. The calibrated $B_i^{(j)}$ are listed in Table 1. In a given experiment, the resonant frequencies will somewhat differ from those

²The intensity $B_i^{(j)}$ is the weighting of the i^{th} cosine term, for the i^{th} resonant frequency of chemical species j . See, e.g., Eq. (3.1).




FIG 1. *Conventional Fourier Transform Spectroscopy for a 50-50 mixture of 2-butanone and cyclohexane. a) Free Induction Decay (FID) Signal. b) Discrete Fourier Transform of FID.*

given in Table 1, depending on, for example, the composition of the mixture in question. In a conventional procedure, one must look for peaks near tabulated reference frequencies, and choose which chemical they are associated with. By contrast, the proposed method in the next sections automatically estimates the resonant frequencies for a given chemical mixture.

Let $y(t)$ be a given FID signal, as shown in Figure 1(a). We Fourier transform $y(t)$ using the Discrete Fourier Transform (DFT) to obtain $\tilde{y}(\omega)$, as in Figure 1(b).³ The conventional procedure is then to search for each peak near an expected resonant frequency. Since the FID undergoes decay, is noisy, etc., and thus does not perfectly conform to the assumptions of a DFT, there will be width about each peak (e.g., the peaks are not delta function spikes, but look more like Gaussian or Cauchy densities), noise, and overlapping peaks which can be hard to differentiate in practice, particularly given that resonant frequencies will shift in a mixture, as seen in Figure 1(b).

Following standard procedure, we identify a peak window for each resonant frequency and sum together the intensity of $\tilde{y}(\omega)$ within this peak window. The summed intensity is then assigned as the total intensity for this peak window. Chemical quantification predictions are sensitive to the peak window, which is chosen heuristically (Malz and Jancke, 2005). Thereafter we are left with K peaks, belonging to chemical j , indexed as $S_k^{(j)}$. We

³As discussed further in section 2, we zero-fill the data $y(t)$ prior to taking the Fourier transform (Malz and Jancke, 2005).

then calculate

$$(2.1) \quad I_j = \sum_k \frac{S_k^{(j)} / B_k^{(j)}}{K},$$

where $B_k^{(j)}$ is the adjusted intensity in Table 1. Using this conventional approach, the concentration of the j^{th} chemical in the mixture is then calculated as

$$(2.2) \quad C_j = I_j / \sum_j I_j.$$

We estimate an error E_j on this estimate C_j using

$$(2.3) \quad E_j = C_j \sqrt{\left(\frac{E_{I_j}}{I_j}\right)^2 + \frac{\sum_j E_{I_j}^2}{(\sum_j I_j)^2}},$$

where

$$(2.4) \quad E_{I_j} = \sqrt{\sum_k \frac{n_k^{(j)} \sigma_S^2}{(K B_k^{(j)})^2}}.$$

$n_k^{(j)}$ is the number of samples along the frequency axis used to defined a peak window for the k^{th} peak of the j^{th} chemical (e.g., the number of samples used to calculate $S_k^{(j)}$), and σ_S is the standard deviation of the noise in the spectral domain $\tilde{y}(s)$.

In short, using a conventional Fourier transform spectroscopy method we calculate the relative concentration of chemical j as C_j in Eq. (2.2), and the uncertainty about this estimate as $2E_j$ in Eq. (2.4) to approximate a 95% credible interval.

Finally, the signal to noise ratio (SNR) of the FID, which we will refer to in the experiments of sections 7 and 8, is given as

$$(2.5) \quad \text{SNR} = \frac{S_k^{(j)}}{B_k^{(j)} \sigma_S}.$$

Using the definition of Eq. (2.5), the SNR will be different for each peak $S_k^{(j)}$. When we quote the SNR, we always state the value for the lowest intensity peak.

In both simulations and real experiments, we consider the FID sampled at 4029 points at $25\mu s$ intervals. For conventional Fourier transform spectroscopy, the time domain data were first apodised with an exponential function and zero-filled to 16384 points to allow a sufficient spectral resolution. Then, a Fourier transform was applied to the FID. A standard baseline correction was then performed on the spectrometer: we fit a polynomial to the

TABLE 1

Resonant frequencies and known intensities for each chemical group in the mixture. The frequencies and theoretical $B_i^{(j)}$ are tabulated from the National Institute of Advanced Industrial Science and Technology Database: <http://sdfs.riodb.aist.go.jp>.

Chemical	Frequencies (ppm)	Excitation profile weighting	Theoretical $B_i^{(j)}$	Calibrated $B_i^{(j)}$
2-Butanone	209.29	90.79%	1	0.908
	36.87	96.71%	1	0.967
	29.43	95.72%	1	0.957
	7.87	93.03%	1	0.93
Cyclohexane	27.1	95.59%	6	5.735

Fourier transform of the FID (the spectrum), and subtract the polynomial fit from the spectrum to product a flat baseline. All such conventional processing was performed in *TopSpin* (Bruker)⁴.

3. Proposed Model Specification. We now propose an alternative to conventional Fourier transform methods for NMR spectroscopy.

Ultimately, we wish to predict the relative concentrations of chemicals in a mixture, from a time-varying free induction decay (FID) signal. Typically the FID is recorded in two channels, $y_1(t)$ and $y_2(t)$, assumed $\pi/2$ out of phase (e.g., real and imaginary parts of a complex signal), and independently corrupted with i.i.d. Gaussian white noise.

Assuming r chemical species in a mixture, we model $y_1(t)$ as

$$\begin{aligned}
 (3.1) \quad y_1(t) = & A_1 \sum_{i=1}^{m_1} B_i^{(1)} \cos((\omega_i^{(1)} - \omega_0)(t + \tau) + \theta) e^{-\alpha t} + \dots \\
 & + A_r \sum_{i=1}^{m_r} B_i^{(r)} \cos((\omega_i^{(r)} - \omega_0)(t + \tau) + \theta) e^{-\alpha t} \\
 & + \epsilon_1(t),
 \end{aligned}$$

with amplitudes A_1, \dots, A_r , frequencies $\{\omega_i^{(j)}\}$ and reference frequency ω_0 in rad/s (chemical species j has m_j resonant frequencies), intensities $\{B_i^{(j)}\}$, decay constant α in 1/s, global phase shift θ in rad, time delay τ in s, and noise $\epsilon_1(t) \sim \mathcal{N}(0, v)$.⁵ We will always assume that the intensities $\{B_i^{(j)}\}$ and reference frequency ω_0 are known and fixed, and we will estimate all other parameters in a procedure outlined in sections 4 and 5. In Figure 2 we show a draw from this generative model, compared with a real FID signal, for a 30% cyclohexane 70% 2-butanone mixture. The real and synthetic FID signals are qualitatively similar – similar resonant frequencies, similar decay rates, etc.

⁴<http://www.bruker.com/products/mr/nmr/nmr-software/software.html>

⁵To convert from the ppm of Table 1 into rad/s, we use $\text{ppm} \times 75 \text{ Hz}$ (for the ^{13}C spectrum on our 7T magnet) $\times 2\pi$.




FIG 2. *a) Real (experimental) and b) Synthetic free induction decay signals for a 30% cyclohexane 70% 2-butanone mixture, sampled at $25\mu\text{s}$ intervals, from a single channel.*

The parameters in the generative model of Eq. (3.1) have clear physical interpretations. The excitation of a sample, e.g. a chemical mixture, in NMR is provided by a radiofrequency (rf) coil which resonates at the frequency of the nucleus of interest, e.g. ^{13}C . The time delay τ in our model arises from the ‘ringdown’ time of this rf coil. The ringdown time is the time taken for the current in the coil to cease. In NMR, rf coils are designed for maximum sensitivity and therefore resonate for a prolonged period after the current supplied to the coil is stopped. The signal from the sample cannot be detected until this ringdown period has finished. In a typical coil this time will be on the order of a few microseconds. In conventional Fourier transform spectroscopy, this delay between excitation and starting the signal detection leads to a linear change in the phase of the signal with frequency variation, which is corrected after Fourier transformation of the raw signal. In our system, we include the time delay in the model directly.

The signal lifetime in NMR is limited by both the relaxation and homogeneity of the magnetic field in which the sample is placed. Both of these effects lead to the decay α of the signal amplitude over time. In an ideal system, the signal will decay through an exponential process (Keeler, 2011). Perhaps surprisingly, this model is also appropriate to a range of non-ideal systems including, e.g., the signal from liquids in rocks (Chen et al., 2005). Here we study pure liquids and so the exponential decay model is most appropriate, though we note that our methodology could easily be adapted for a variety of other models describing the signal decay.

We can rewrite Eq. (3.1) as an inner product of amplitude coefficients $\mathbf{a} = (A_1, \dots, A_r)^\top$ with cosine basis functions $\phi = (\phi_1, \dots, \phi_r)$ parametrized by $\psi = \{\{\omega\}_i^j, \theta, \tau, \alpha\}$:

$$(3.2) \quad y_1(t) = \mathbf{a}^\top \phi(t, \psi) + \epsilon_1(t),$$

$$(3.3) \quad \epsilon_1(t) \sim \mathcal{N}(0, v).$$

The model for $y_2(t)$ is the same,

$$(3.4) \quad y_2(t) = \mathbf{a}^\top \boldsymbol{\varphi}(t, \psi) + \epsilon_2(t),$$

$$(3.5) \quad \epsilon_2(t) \sim \mathcal{N}(0, v),$$

except we use sine instead of cosine basis functions in $\boldsymbol{\varphi}$ to account for the $\pi/2$ phase difference between the two channels.

In section 4 we infer $p(\mathbf{a}|\mathcal{D}, \{t_n\}_{n=1}^N, \psi, v)$, a posterior distribution over the amplitude variables given an observed FID signal \mathcal{D} at times $\{t_n\}_{n=1}^N$, and nuisance variables ψ and v . Since the relative concentrations of chemicals i and j is given by $|A_i/A_j|$, this posterior distribution can be used to estimate, with uncertainty intervals, the relative concentrations of chemicals in a mixture.

In section 5 we discuss estimation of frequency variables $\{\omega\}_i^{(j)}$, other nuisance variables in ψ , and the noise variance v . Performance in estimating chemical concentrations is especially sensitive to frequency estimates. Moreover, accurately estimating frequency parameters is difficult because the likelihood of the data is highly multimodal as a function of these frequency variables.

Incidentally, it may be a slight misnomer to refer to ψ, v as nuisance variables. Although we are primarily interested in estimating the relative concentrations of chemicals in a mixture – and thus, the \mathbf{a} variable – the learned frequencies, decay rates, phase shifts, and noise variance are all still enlightening. For instance, the learned frequencies could be used, in principle, to better understand how resonant frequencies shift for particular chemicals when immersed in a mixture.

4. Inference over Amplitudes. We wish to infer $p(\mathbf{a}|\mathbf{y}_1, \mathbf{y}_2, \psi, v)$, where \mathbf{y}_1 and \mathbf{y}_2 are the “real” and “imaginary” channels respectively observed at times $\{t_{1n}\}_{n=1}^N$ and $\{t_{2m}\}_{m=1}^M$. To reflect our prior uncertainty over the amplitudes, we specify a prior $p(\mathbf{a}) = \mathcal{N}(\boldsymbol{\mu}_0, S_0)$, and find

$$(4.1) \quad p(\mathbf{a}|\mathbf{y}_1, \mathbf{y}_2, \psi, v) = \mathcal{N}(\boldsymbol{\mu}, S),$$

$$(4.2) \quad S = (S_0^{-1} + \frac{1}{v}\Phi^\top\Phi + \frac{1}{v}\Psi^\top\Psi)^{-1},$$

$$(4.3) \quad \boldsymbol{\mu} = S(S_0^{-1}\boldsymbol{\mu}_0 + \frac{1}{v}\Phi^\top\mathbf{y}_1 + \frac{1}{v}\Psi^\top\mathbf{y}_2),$$

where Φ and Ψ are respectively $N \times r$ and $M \times r$ matrices with entries

$$(4.4) \quad \Phi_{i,j}(\psi, t_i) = \phi_j(\psi, t_{1i}),$$

$$(4.5) \quad \Psi_{i,j}(\psi, t_i) = \varphi_j(\psi, t_{2i}).$$

We may not know anything about the relative concentrations of the chemicals a priori, in which case we wish to have a vague uninformative prior over the amplitudes. For example, one can use an improper uniform prior on the amplitudes by taking $S_0 = \lim_{\gamma \rightarrow \infty} \gamma I$,

which causes the S_0^{-1} and $S_0^{-1}\mu_0$ terms to vanish in Equations (4.2) and (4.3), respectively. Alternatively, one can start with a vague (much higher variance than mean) prior (effectively a uniform prior), and then learn the value of γ from the data following the procedure outlined in section 5.

The posterior distribution over the amplitudes is conditioned on the nuisance variables ψ and v . Before we can use this posterior to make predictions about the relative concentrations of chemicals in a mixture, we must first either integrate away these nuisance variables, or find a point estimate, conditioned on the data.

5. Estimating Nuisance Variables. We find that the performance of the model in Equations (3.3) and (3.4) is sensitive to estimates of the nuisance variables, particularly the frequency variables. Furthermore, estimating these frequencies is non-trivial, due to a severely multimodal likelihood.

As a step towards estimating the nuisance variables, ψ and v , we can analytically integrate away the amplitude variables from the likelihood of the data $p(\mathbf{y}_1, \mathbf{y}_2 | \mathbf{a}, \psi, v)$:

$$(5.1) \quad p(\mathbf{y}_1, \mathbf{y}_2 | \psi, v) = \int p(\mathbf{y}_1, \mathbf{y}_2, \mathbf{a} | \psi, v) d\mathbf{a}$$

$$(5.2) \quad = \int p(\mathbf{y}_1, \mathbf{y}_2 | \psi, v, \mathbf{a}) p(\mathbf{a}) d\mathbf{a}$$

$$(5.3) \quad = \int p(\mathbf{y}_1 | \mathbf{a}^\top \boldsymbol{\phi}(t, \psi), v) p(\mathbf{y}_2 | \mathbf{a}^\top \boldsymbol{\varphi}(t, \psi), v) p(\mathbf{a}) d\mathbf{a}$$

$$(5.4) \quad = \int p(\mathbf{a} | \mathbf{y}_1, \psi, v) p(\mathbf{y}_1 | \psi, v) p(\mathbf{y}_2 | \mathbf{a}^\top \boldsymbol{\varphi}(t, \psi), v) d\mathbf{a}.$$

In Equation (5.3), the likelihood of the data decomposes into a product of Gaussian likelihoods for each channel, since the channels are independent given the noise free signal. Performing the integration in Eq. (5.4), assuming $S_0 = \gamma I$ for notational simplicity, we find the log *marginal* likelihood of the data is

$$(5.5) \quad \begin{aligned} \log p(\mathbf{y}_1, \mathbf{y}_2 | \psi, v) &= -\frac{1}{2v} (\|\mathbf{y}_1 - \Phi \mu_{\text{real}}\|^2 + \|\mathbf{y}_2 - \Psi \mu\|^2) - \frac{1}{2\gamma} \|\mu - \mu_0\|^2 \\ &- \frac{1}{2} (\mu - \mu_{\text{real}})^\top \left(\frac{\Phi^\top \Phi}{v} + \frac{I}{\gamma} \right) (\mu - \mu_{\text{real}}) - \frac{N+M}{2} \log(v) - \frac{r}{2} \log(\gamma) - \frac{N+M}{2} \log(2\pi), \end{aligned}$$

where μ_{real} is the posterior mean on the amplitudes if we were to only have data from the real channel:

$$(5.6) \quad \mu_{\text{real}} = \left(\frac{\Phi^\top \Phi}{v} + \frac{I}{\gamma} \right)^{-1} \left(\frac{\mu_0}{\gamma} + \Phi^\top \mathbf{y}_1 \right).$$

The first two terms of Eq. (5.5) are model fit terms, since $\Phi \mu_{\text{real}}$ and $\Psi \mu$ are estimates of the real and imaginary channels of the noise free FID. The remaining terms are normalization

constants, and automatically calibrated complexity penalties (Rasmussen and Ghahramani, 2001), which come from integrating away the amplitudes from the likelihood.

We can estimate the nuisance variables ψ, v by finding the $\hat{\psi}, \hat{v}$ that maximize the log marginal likelihood $\log p(\mathbf{y}_1, \mathbf{y}_2 | \psi, v)$ in Eq. (5.5), a procedure sometimes called *empirical Bayes* or *type-II maximum likelihood*, or we can use the marginal likelihood to integrate away ψ, v via Markov chain Monte Carlo (MCMC), in order to sample from the posterior $p(\mathbf{a} | \mathbf{y}_1, \mathbf{y}_2)$. In either case, we must navigate a challenging multimodal likelihood surface.

In Figure 3, we illustrate the behaviour of the marginal likelihood in Eq. (5.5) as a function of the nuisance parameters ψ, v , on a synthetic equal mixture of cyclohexane and 2-butanone. We generate the data by sampling from the generative model in Eqs. (3.3)–(3.4) with frequencies and intensities for cyclohexane and 2-butanone as given in Table 1.

Each panel in Figure 3 shows how the marginal likelihood varies as a function of a single parameter in the model of Eq. (3.1), with all other parameters set to the true generating values. The true values are shown with dashed vertical lines. Even in simulation, the likelihood surface is severely multimodal as a function of the frequency parameters. Moreover, there are strong dependencies between the frequency and phase parameters. Therefore exploring this likelihood surface is non-trivial. Our prior knowledge of the frequencies is only accurate to within about 1000 rad/s (or 2 ppm). We can see in Figure 3(a)–(f) that this range spans multiple local optima. A naive gradient based optimizer – greedily choosing to improve the value of an objective function on every iteration – will converge to undesirable local optima, rendering predictions of the amplitudes unusable. Likewise, popular sampling schemes, such as Metropolis Hastings, become stuck in a local optimum if using a small step size, and proposals with a large step size are rejected. We can also see that sampling may have limited value in this application, since the global optima in the frequency space are sharply peaked. Moreover, while the posterior over the amplitude parameters is sensitive to the estimated values of the frequency parameters, this posterior is not so sensitive to noise variance or decay that exploring the unimodal distributions shown in Figure 3 would provide major performance gains.

The local optima arise because a signal composed of a sum of frequency components can be extremely well modelled in parts, while missing the rest of the signal entirely, by a range of erroneous frequencies. The smaller the signal, the less pronounced the multimodality. Increasing an estimate on the initial decay rate α decreases the amount of available information (since the signal is then assumed to decay away quickly), and therefore decreases the multimodality in the likelihood. Thus increasing the estimate of α has an annealing effect on the likelihood surface. Increasing the noise variance has a similar effect: the more noise, the less available signal. Indeed there are natural dependencies between the decay rate and noise variance parameters. An overly high decay rate means a large portion of the signal can only be explained as noise. In short, systematically overestimating the decay rate, as part of a gradient based optimization procedure, can be used to help locate a global optimum. However, we find an explicit simulated annealing simplex algorithm (SIMPSA) (Cardoso, Salcedo and Feyo de Azevedo, 1996; Nelder and Mead, 1965), which was espe-




FIG 3. Log marginal likelihood plots as a function of all the nuisance parameters in simulation. Each panel shows how the log marginal likelihood varies with a given variable, with all other variables set to their true values. The ground truth frequencies are 206.29, 36.57, 28.43, 7.77 ppm for butanone and 27.1 ppm for cyclohexane. The dashed line gives the corresponding true value for the parameter of interest. The FID data were generated from 2-butanone (30%) and cyclohexane (70%) mixture with 8 scans, corresponding to an SNR (as defined in Eq. (2.5) of section 2) of 26 for 2-butanone and 381 for cyclohexane).

cially designed for continuous, bounded, global optimization, more reliable for this NMR spectroscopy application.

In short, to estimate relative chemical concentrations, we follow a three step procedure:

1. Use SIMPSA to find the $\psi_{\text{est}}, v_{\text{est}}$ that optimize the marginal likelihood $\log p(\mathbf{y}_1, \mathbf{y}_2 | \psi, v)$ in Equation (5.5), using a vague uninformative prior $p(\mathbf{a})$ on the amplitudes.
2. Analytically infer the posterior $p(\mathbf{a} | \mathbf{y}_1, \mathbf{y}_2, \psi_{\text{est}}, v_{\text{est}})$ as in Equation (4.1), conditioned on the maximum marginal likelihood estimates of the nuisance parameters ψ, v .
3. By sampling from $p(\mathbf{a} | \mathbf{y}_1, \mathbf{y}_2, \psi_{\text{est}}, v_{\text{est}})$, sample from the distribution over the relative concentrations $r_{ij} = |A_i/A_j|$ of chemicals i and j , $p(r_{ij} | \mathbf{y}_1, \mathbf{y}_2, \psi_{\text{est}}, v_{\text{est}})$, where A_i is the i^{th} component of \mathbf{a} .

The computational demands of the proposed procedure are dominated by the need to evaluate the marginal likelihood of Eq. (5.5) for a variety of settings of ψ and v . Each such evaluation costs $\mathcal{O}((N+M)r^2+r^3)$ operations for N and M points respectively in the real and imaginary channels, and r chemical species: it takes $(N+M)r^2$ operations to compute $\Phi^T \Phi + \Psi^T \Psi$, and $\mathcal{O}(r^3)$ operations to take the Cholesky decomposition of this term (for solving linear systems). Thus for a system with a fixed number of chemicals, the required computational operations scale linearly with the number of collected datapoints.

6. Related Work. There is a body of groundbreaking work using models motivated for NMR spectroscopy, similar in form to Equation (3.1). In this section we briefly describe a selection of this work, and some limitations. Yoon et al. (2006) provide a general review of statistical models with various NMR applications.

Bretthorst (1990) was an early pioneer of Bayesian estimation of a quadrature model, similar to Eq. (3.1), but without the local phase shift τ . While inference over the amplitude variables follows a similar procedure to that described in section 4, inference over the frequency variables involves several approximations, including a quadratic Taylor series expansion, leading to a Gaussian posterior over frequencies. However, in general, the posterior over frequencies can be highly non-Gaussian (as shown in Figure 3). Generally, these approximations are not robust to a multimodal likelihood surface, and tend to break down when there are more than a few resonant frequencies to estimate. Evilia, Effiong and Whittenburg (1993) contains a similar early quadrature NMR model, but the exact model specification and estimation procedures are unclear, and the model is tested on systems with at most 2 resonant frequencies.

Andrec and Prestegard (1998) use Metropolis Hastings to sample from the posterior distribution over parameters in a quadrature model similar to Bretthorst (1990), ultimately to estimate coupling constants in antiphase doublets. Andrec and Prestegard (1998) do not estimate more than two frequency parameters at a time, and do not compare with conventional spectroscopy. Furthermore, Metropolis Hastings (MH) is not generally suitable for exploring multimodal likelihood surfaces, as discussed in section 5. In short, using an MH proposal distribution with a small width will cause the sampler to become trapped

in undesirable local optima, while a large width proposal is extremely unlikely to find a global optimum, and if it does, it will never move from a point estimate, which defeats the purpose of sampling.

Dou and Hodgson (1995) perform Gibbs sampling (Geman and Geman, 1984) for parameter estimation, with similar performance to Bretthorst (1990). Gibbs sampling involves alternately sampling one parameter while conditioning on the others, in a cycle. When there are strong dependencies between parameters, such as frequency and phase parameters, and multimodality, Gibbs sampling is known to mix poorly (?). Therefore a Gibbs sampling procedure may struggle in many NMR applications.

Rubtsov and Griffin (2007) focus on estimating the number of components (e.g., resonant frequencies) in a model similar to Bretthorst (1990), using reversible jump Markov chain Monte Carlo (MCMC). While the model is applied to experimentally acquired data, the experiments are taken in the ^1H spectrum, where the signal to noise ratio is much higher than in the ^{13}C spectrum.

In recent work, Hutton et al. (2009) develop a specialised model for mixtures of water and metabolites. The model has a similar form to Bretthorst (1990), except the amplitude coefficient for water is time-varying. Estimation of frequency parameters follows a simulated annealing MCMC approach. It is difficult to determine whether the methodology in Hutton et al. (2009) is generally applicable to quadrature NMR models, since the paper is focused on a specialised application in the (^1H) proton spectrum, where the SNR is much higher than in the carbon spectrum (considered in our paper), there is no comparison to conventional spectroscopy, and there is no motivation or description of the applied sampling scheme, except that the computational costs can be considerable.

Rather than construct a model in the time domain, Astle et al. (2012) consider modelling preprocessed data (apodised, phase corrected, and with a baseline correction) in the frequency domain, leveraging positivity requirements in this domain. While promising, the model relies on conventional preprocessing and implicit assumptions of the discrete Fourier transform (DFT), and thus may not be as general as time domain models.

Furthermore, Rubtsov et al. (2010) develop a promising statistical procedure for estimating resonant frequencies in the ^1H spectrum, independently from quadrature NMR models, and compare to binning. And Aboutanios, Kopsinis and Rubtsov (2012) develop heuristics for estimating components in a quadrature NMR model, which show some robustness to synthetic noise; however, the dataset considered is synthetic and has a well defined discrete Fourier transform.

Our paper focuses on aspects of quadrature NMR modelling which have not been thoroughly explored in prior work: 1) we develop a statistical model for NMR spectroscopy with a novel chemical quantification application, 2) we discuss challenges in estimating quadrature NMR models, such as a multimodality, and propose solutions, 3) we provide thorough quantitative comparisons with conventional Fourier transform spectroscopy, and 4) we consider both synthetic and experimentally acquired data in the ^{13}C spectrum, with low SNR, and many resonant frequencies.

7. Simulations: 2-Butanone Cyclohexane Mixture. To better understand the behaviour of the proposed model (sections 3-5), and conventional Fourier transform (FT) spectroscopy (section 2), we first predict the relative concentrations of 2-butanone and cyclohexane in mixtures from synthetic free induction decay (FID) signals – a controlled environment with a ground truth.⁶ We simulate FID signals using the generative model of Equation (3.1), using the theoretical intensities from Table 1. Note that in real experiments, resonant frequencies will shift depending on the composition of a mixture, especially in the presence of commonly used solvents such as CDCl_3 . Thus to make the simulations more representative of real experiments, and more challenging, we let the ground truth frequencies differ from the tabulated values by ± 3 ppm, and initialise the estimation procedure of section 5 using the values in Table 1.

As can be seen in Figure 2, a synthetic FID from Eq. (3.1) resembles an actual FID response to a 2-butanone cyclohexane mixture. In section 8 we then compare these models on experimentally acquired FID signals.

7.1. Stress Tests. We start by stress testing the proposed model in response to varying signal to noise ratios (SNR) (as defined in Eq. (2.5)) in the FID signal.

Taking a mixture of 30% 2-butanone 70% cyclohexane as an example, estimates and uncertainties regarding the concentration of 2-butanone, using the proposed model, are shown in Figure 4, with the true concentration of 2-butanone (30%) given by the dashed line. The gray shade indicates a concentration between 27%–33%. The error bars indicate a 95% credible interval (two sample standard deviations of $p(|A_i/A_j||\mathbf{y}_1, \mathbf{y}_2, \psi, v)$ using 10000 samples and the proposed model of section 3).

To test the robustness of the proposed model to noise, we selected a set of SNR values, and generated data 30 times for each of these values. As shown in Figure 4, above a SNR of 4, the estimated concentrations are located within the gray shade (27%–33% 2-butanone), and contain the true concentration (30%) in 29/30 experiments. As expected, the estimated concentrations decrease in variance across different datasets as SNR increases, and converge to the true generating values. Figure 4 also shows that below an SNR of around 4, the proposed model tends to overestimate the concentration of 2-butanone.

We now analyse the effect of increased noise on the marginal likelihood of the data, while keeping all other parameters fixed. Since the four resonant frequencies of 2-butanone behave similarly, we analyse the frequency near 206.29 ppm.

Figure 5 displays the log marginal likelihood as a function of the frequency. The SNR gradually decreases from (a) to (d), with 5.9 for (a), 4.2 for (b), 3.0 for (c), and 2.1 for (d). As shown in Figure 5, the true frequency (indicated by the dashed line) becomes less differentiated as the SNR decreases, and the likelihood surface becomes increasingly multimodal, with no global optimum near the true generating frequencies. The precision of

⁶For conventional spectroscopy, we use the known true phase shifts to perform an exact phase correction. On the other hand, the proposed Bayesian model of sections 3-5 automatically learns phase corrections from the data.




FIG 4. The relative concentration (and uncertainty) of 2-butanone in synthetic cyclohexane 2-butanone mixtures at various SNR, using the proposed model. The true concentration (30%) is given by the dashed line, with the shaded area representing 27%–33% 2-butanone. The error bars represent the 95% credible region.

estimating the frequencies as well as the final concentrations at selected SNR were analysed by 30 repetitions and are given in Table 2.

In Figure 5(a) the true frequency is well-resolved, and indeed the proposed approach estimates the frequencies accurately, as given in Table 2. As the SNR decreases to 4.2 (Figure 5(b)), the marginal likelihood profile becomes less peaked at the true frequency. At such an SNR, the credible sets for the proposed method become relatively large, as shown in Table 2, and the method is more susceptible to the local optima shown in the likelihood surface of Figure 3. At this SNR level, the issue of local optima is largely alleviated with multiple SIMPSA restarts.

Figures 5(c) and 5(d) show the marginal likelihood at SNR values of 3.0 and 2.1 re-

TABLE 2
Bayesian estimates of the 2-butanone (30%) and cyclohexane (70%) mixture.

SNR	Error in frequency estimation (ppm) (mean \pm 1 std. dev.)	Bayesian 95% credible interval	Empirical Coverage with 68% and (95%) credible intervals	Percentage of reconstructions within 27%-33% 2-butanone
9.3	-0.0039 ± 0.0096	1.09%	91% (100%)	100%
5.9	-0.0072 ± 0.0043	3.44%	50% (100%)	100%
4.2	-0.0067 ± 0.0187	4.92%	82% (100%)	82%
3.0	0.3383 ± 0.7159	6.80%	64% (73%)	60%
2.1	0.5608 ± 1.6385	8.34%	27% (82%)	9%




FIG 5. Log marginal likelihood plots as a function of a resonant frequency for butanone, with the true frequency (206.29 ppm) indicated by the dashed line. The corresponding SNR for each panel are (a) 5.9, (b) 4.2, (c) 3.0, and (d) 2.1. For high SNR values, (a) and (b), the peak at the true frequency is clearly identifiable.




FIG 6. Comparison of the Bayesian and FT performances on the same FID datasets. 30 repetitions are generated for 30% 2-butanone at an SNR of 4.2. The dashed line indicates the true concentration and the error bars correspond to 95% credible regions.

spectively. As the SNR increases, predictions with the proposed Bayesian model becomes increasingly accurate, converging to the truth, as shown in the second and last column in Table 2. When the SNR is above 4.2, the proposed Bayesian method’s 95% posterior credible interval always contains the true concentration.

7.2. *Simulated Comparison to Conventional Spectroscopy: High Concentrations.* We now examine the behaviour of both models in predicting high concentrations of a given chemical.

Results on a sample of 30% 2-butanone at an SNR of 4.2 are shown in Figure 6, with the true concentration indicated by the dashed line and the 95% credible region given by the error bar. To confirm reproducibility, datasets were generated 30 times with the same concentration and SNR. Figure 6 shows that the proposed approach consistently makes accurate predictions that always include the true value in the 95% credible interval. Its estimates are always within a $30\% \pm 3.5\%$ bound.

Conversely, predictions using the conventional FT approach vary significantly in a $\pm 20\%$ concentration interval, and even with broader error bars than the proposed approach FT estimates are typically unable to cover the true concentration. The FT approach is particularly sensitive to noise, which will become even more apparent when we study low concentrations.

We perform a similar analysis for 2-butanone concentrations ranging from 10% to 90%. For clarity, representative results from only two datasets at each concentration were plotted

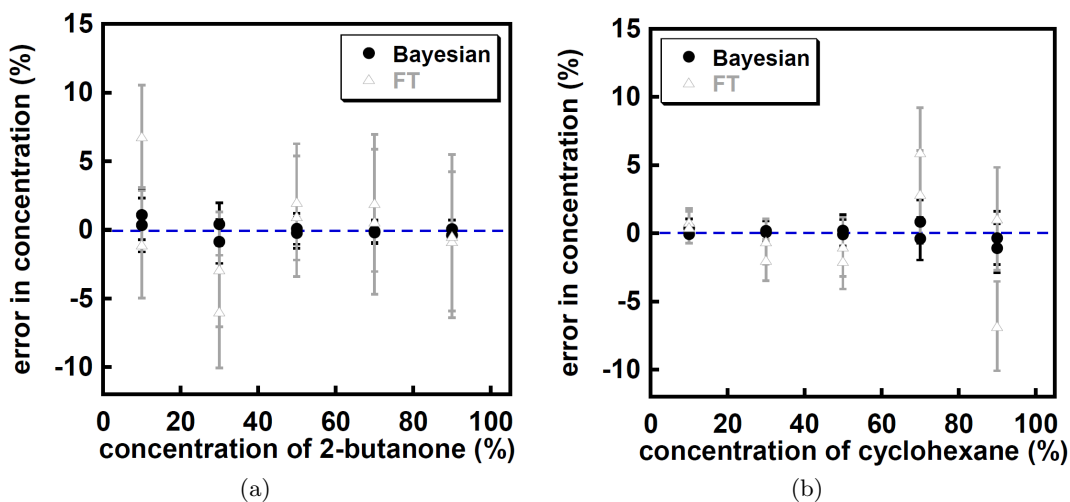


FIG 7. The proposed Bayesian and conventional FT estimations for varying relative concentrations of 2-butanone and cyclohexane in simulation. The noise levels for all the data are the same and are chosen to allow the SNR to be as low as possible but still retain acceptable performance. In particular, the SNR is 4.5 for the mixture containing 10% 2-butanone. The vertical axis shows absolute errors. The proposed Bayesian approach is more accurate and robust than the conventional FT approach in terms of both expected mean and uncertainty.

in Figure 7, with the same noise level.⁷

The Bayesian uncertainty intervals generally decrease as the concentration of 2-butanone rises from 10% to 90%. Notice in Table 1 that cyclohexane has a single intensity near 6, whereas 2-butanone has 4 intensities near 1. Due to the lower intensities of 2-butanone, relative to cyclohexane, it is more difficult to differentiate the signal from 2-butanone from the noise, and so predictions of both cyclohexane and butanone concentrations ought to be more accurate and certain as the relative concentration of 2-butanone rises. Similarly, uncertainty in the FT predictions of cyclohexane increase with relative increases in cyclohexane concentrations (decreases in 2-butanone).

In the cases presented, the proposed Bayesian approach consistently produces more accurate point predictions and smaller and more accurate uncertainty intervals than a conventional Fourier transform approach.

7.3. Simulated Comparison to Conventional Spectroscopy: Low Concentrations. We now investigate the performance of the proposed Bayesian approach compared to the conventional FT approach on low concentrations ($< 10\%$). Low concentrations are of particular interest, as spectroscopists often wish to determine the presence or absence of a chemical in a

⁷When we refer to *absolute error*, we mean the absolute distance between the prediction and the true concentration percentage, and by *relative error* we mean the ratio of the absolute error to the true concentration percentage.

system, or whether a particular chemical falls under a given concentration threshold. As we will see, the lower the concentration of a chemical, the harder the concentration is to accurately estimate.

We begin with a simulation using 5% 2-butanone with an SNR of 6. Figure 8(a) shows the whole FT spectrum for this dataset, with the locations of expected peaks indicated by the dashed lines. It is very difficult to see the low intensity peaks from 2-butanone. Figure 8(b) zooms in to part of the spectrum indicated by the red box in 8(a). As shown in Figure 8(b), a peak from the dilute species, 2-butanone, overlaps with a high intensity peak from the concentrated species cyclohexane, which makes assignment of this peak difficult, causing the concentration of dilute chemicals to be overestimated. Moreover, the duration of the acquisition is limited owing to the radio frequency (rf) pulse used for proton decoupling (Keeler, 2011). This limited acquisition time causes truncation artifacts on the cyclohexane peak at 27.3 ppm, making it more difficult to differentiate the butanone peak at 28.4 ppm.

The Bayesian log marginal likelihood as a function of the same frequency from 2-butanone is displayed in Figure 8(c), with the same frequency window as used for the FT in 8(b).

As shown in Figure 8(c), using the proposed Bayesian method, true resonant frequencies of dilute species are well-resolved even in the presence of a nearby intense peak, due to the principled models of decay and noise in the FID signal. Such log marginal likelihood shape reduces the difficulty of learning the right frequency in the Bayesian procedure and, as shown in Figure 8(d), the final probability distribution is well resolved with a 2.4% relative error (0.12% absolute error).

Simulated results for 3% to 7% 2-butanone are shown in Figure 9, with the true concentration given by the dashed line. The FT predictions are more variable than the Bayesian estimates, and systematically overestimate the concentration in almost all cases. Conversely, the Bayesian results have no noticeable bias and the true concentrations are always within the credible region. The smaller error bar (0.7%) – which has been reliable in these experiments – also allows the Bayesian approach to reliably distinguish between small concentration differences.

In conclusion, the analysis of the synthetic data at low concentrations ($< 10\%$) demonstrates the strength of the Bayesian approach: it is more accurate than the FT, particularly in peak overlapping situations, or when the SNR or chemical concentration levels are particularly low.

8. Experiments. We now compare the proposed Bayesian model with a conventional spectroscopy approach on experimentally acquired FID signals corresponding to mixtures of cyclohexane and 2-butanone.

Cyclohexane ($\geq 99.5\%$) and 2-butanone ($\geq 99.5\%$) were purchased from Sigma Aldrich. The mixtures were prepared by mixing the two chemicals with various concentrations based on weight. To confirm reproducibility of results, each mixture was made twice and both preparations were used in the experiments. The error in sample preparations was calculated

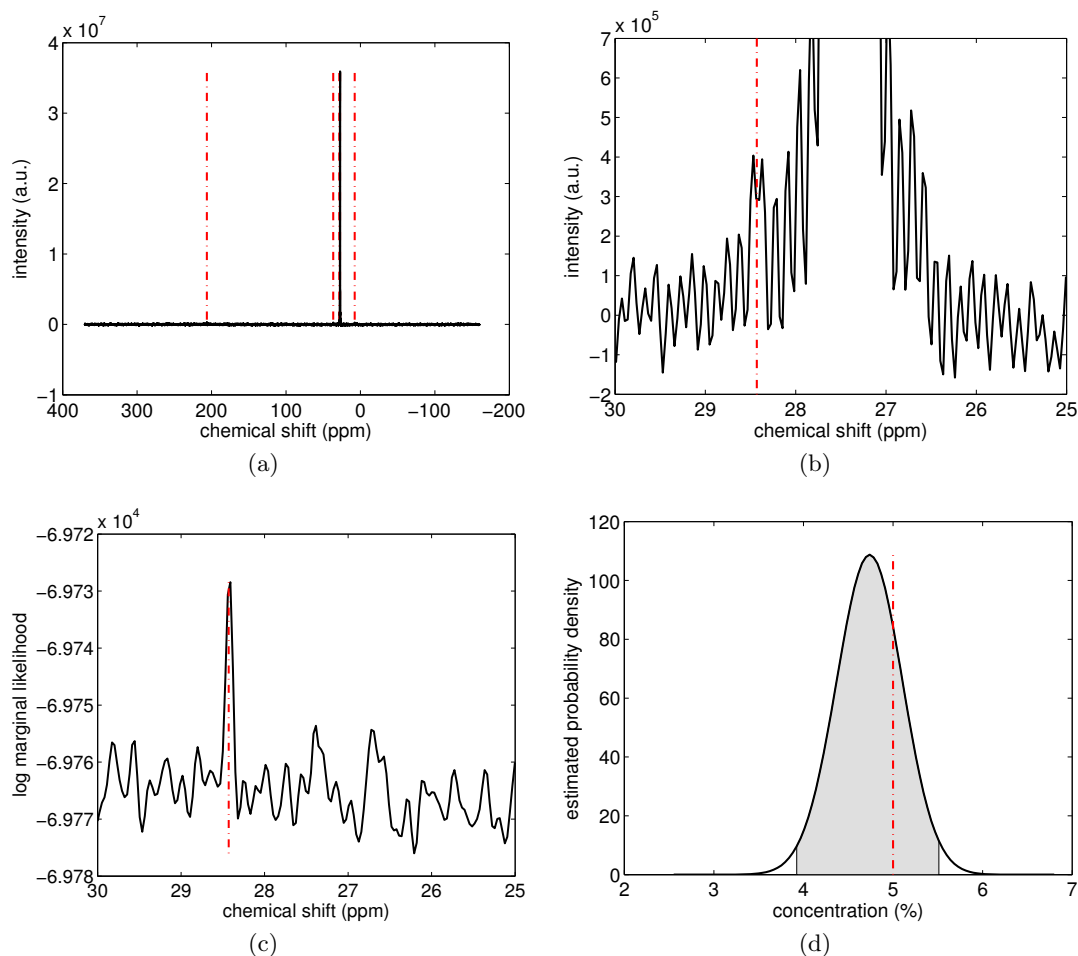


FIG 8. Comparison between the proposed Bayesian model and conventional Fourier transform spectroscopy on the same synthetic FID dataset. The ground truth frequencies are 206.29, 36.57, 28.43, 7.77 ppm for butanone and 27.3 ppm for cyclohexane. The data were generated assuming a mixture of 5 percent 2-butanone with an SNR of 5.9 (a) The whole FT spectrum. (b) Zoom in to part of the spectrum indicated by the red box in (a). The peak (28.43 ppm) from 2-butanone is on the shoulder of the peak (27.3 ppm) from cyclohexane, and is prone to overestimation owing to the truncation artifacts arising from that peak. (c) Log marginal likelihood as a function of the 28.43 ppm resonance frequency. It is well distinguished from adjacent frequencies. (d) the concentration distribution of 2-butanone estimated from the proposed Bayesian model. The dashed lines indicate the true values. As shown in (c) truncation artifacts do not interfere with the performance of the proposed Bayesian model.




FIG 9. Simulated comparison of the Bayesian and FT approaches at low concentrations of 2-butanone. FID data were simulated with the a minimum noise level for acceptable performance (an SNR of 4.4, which is the smallest SNR at which the 3% concentration can be reasonably estimated). The true concentrations are given by the dashed line.


from the uncertainty in the electronic scale.

Conventional spectroscopy experiments were performed on a Bruker DMX300 spectrometer operating at a ^{13}C resonance frequency of 75.47 MHz with a 5mm microimaging probe. The reference frequency ω_0 was calibrated using the standard reference tetramethylsilane (TMS).

The default digital filter on the spectrometer performs an unknown transformation (for proprietary reasons) on the FID signal, making the generative model in Eq. (3.1) inapplicable. We therefore acquire the FID signal digitally, to minimise the noise in acquisition, and then use the TopSpin (Bruker) software on the spectrometer to convert the signal to a “pseudo-analogue” form. This approach maintains the integrity of the generative model in Eq. (3.1). We found the performance of the conventional Fourier transform method is equivalent using the digital or pseudo-analogue form of the FID signal.

Experimentally, the time domain FID data were acquired using a single 90° pulse on ^{13}C channel with broadband ^1H decoupling, as shown in Figure 10. Each FID data set consists of 4096 complex pairs with a $25\mu\text{s}$ sampling interval. The recycle time (the time between excitations of the system) was set as 180s to allow sufficient relaxation recovery. The digitally filtered FID data were converted to pseudo-analogue data format after acquisition using TopSpin software, which modified the number of data points to 4029.

Phase correction was manually performed on the spectrometer using both zero and first order phase corrections (θ and τ in Eq. (3.1)) to ensure all peaks are in the so-called “absorption mode” (Keeler, 2011).

FIG 10. *NMR on-pulse sequence.*

Because the NMR spectrometer can only excite the spins within a certain resonance frequency range, the extent to which each chemical group is excited differs from each other slightly. To account for this behaviour, we adjust the intensities $B_i^{(k)}$ in the signal model of Eq. (3.1) as described in section 2 and Table 1.

8.1. *Experimental Comparison to Conventional Spectroscopy: High Concentrations.* We first perform experimental measurements on mixtures of reasonably high concentrations, from 10% to 90%. Similar to the simulations, the noise level was set by fixing the number of scans in experiments, with an SNR for the lowest concentration just above a threshold of 4.5, determined in section 7.2. To confirm the ground truth of the concentration, data were acquired from two prepared samples at each concentration. Estimation errors are shown in Figure 11. Similar accuracy and general behaviour as in the simulations (Figure 7) are observed. Accordingly, the proposed Bayesian approach has far less error and more realistic uncertainty estimates than the conventional Fourier transform approach.

8.2. *Experimental Comparison to Conventional Spectroscopy: Low Concentrations.* We now investigate experimentally the ability of the proposed method and conventional spectroscopy to estimate low concentration levels of 2-butanone. We prepare sets of samples with concentrations ranging from 3% to 7% in 1% increments. We create two sets per concentration level and perform 3 independent measurements on each sample. The SNR for the 3% cases is 4.4. The Bayesian and FT results are shown in Figure 12, with the expected concentration (ground truth determined during mixture preparation) indicated by the dashed line. The error bars along the horizontal direction represent the error in sample preparation. The FT method consistently overestimates the concentration whereas there is no systematic bias from the Bayesian results. Moreover, the FT results from the

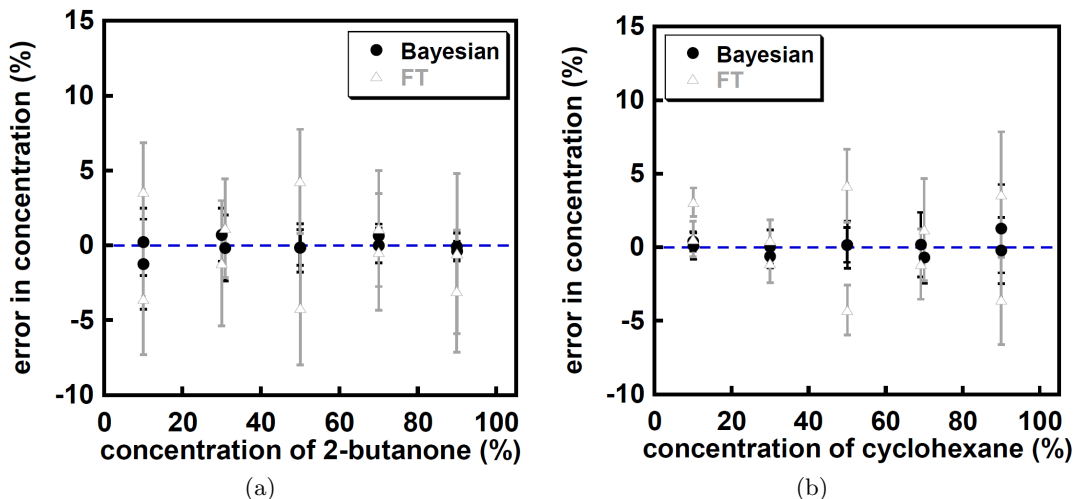


FIG 11. *Experimental results from Bayesian and FT approaches. The concentrations and SNR were set similar to the simulation cases shown in Figure 7. To avoid bias in the ground truth, data were acquired on two sets of samples. Error bars along the x axis were plotted but are too small to be seen.*

same sample can scatter over a 5% range and their error bars are unrealistically small. In contrast, the Bayesian predictions are always within a 1% concentration bound and with more reliable uncertainty estimates.

A particular set of data from the low concentration cases were taken to provide a more visual description. Figure 13 shows the results of a dataset from a 3% 2-butanone mixture. The conventional FT spectrum, Figure 13(a), does not show the peaks from the 3% 2-butanone. It is difficult to differentiate most of the peaks in 2-butanone, as shown in Figure 13(b). However, the proposed Bayesian approach can still infer the true concentration with reasonable accuracy and uncertainty.

9. Discussion. We have proposed an alternative method for nuclear magnetic resonance spectroscopy. Unlike conventional Fourier transform spectroscopy, the proposed approach explicitly models decay, noise, frequencies, intensities, and phase shifts, and can leverage prior information in a principled way using probability distributions. We compared the proposed Bayesian approach with conventional Fourier transform spectroscopy in simulations and on experimentally acquired free induction decay signals, for quantifying relative chemical concentrations in a mixture. The proposed approach significantly outperformed conventional spectroscopy, with more accurate estimates and uncertainty intervals, particularly in low signal to noise ratio cases, and in cases where there were overlapping peaks in the Fourier transform of the free induction decay. In summary, the proposed method – while sensitive to frequency estimation – can be used with little human intervention, for reproducible and accurate estimates of chemical concentrations, and may be used in




FIG 12. *Experimental comparison between the Bayesian model and FT for inferring low concentration levels. The data were acquired at the same noise level that allows the lowest concentration to be robustly estimated. Two samples were prepared at each concentration with three independent datasets acquired from each sample.*




FIG 13. *Experimental findings on a $3.07\% \pm 0.20\%$ 2-butanone mixture. (a) Conventional Fourier spectrum. (b) Zoom in to the low intensity 2-butanone peaks. The peaks are almost at the same height as the noise level and cannot be assigned and quantified accurately. (c) Concentration probability distribution of 2-butanone calculated from the Bayesian model. The shaded area gives the 95% credible region.*

detecting the presence or absence of chemicals where conventional Fourier transform spectroscopy cannot be used. In general, the robustness of the proposed method to low signal to noise ratios, and overlapping spectral peaks, has widespread promise in analytic chemistry.

Acknowledgements. The authors would like to thank NSERC, Microsoft Research Connections, and the EPSRC (Grants No. EP/F047991/1 and EP/K039318/1) for financial support.

References.

- ABOUTANIOS, E., KOPSINIS, Y. and RUBTSOV, D. (2012). Instantaneous frequency based spectral analysis of nuclear magnetic resonance spectroscopy data. *Computers & Electrical Engineering* **38** 52–67.
- ANDREC, M. and PRESTEGARD, J. H. (1998). A Metropolis Monte Carlo implementation of Bayesian time-domain parameter estimation: application to coupling constant estimation from antiphase multiplets. *Journal of Magnetic Resonance* **130** 217–232.
- ASTLE, W., DE IORIO, M., RICHARDSON, S., STEPHENS, D. and EBBELS, T. (2012). A Bayesian model of NMR spectra for the deconvolution and quantification of metabolites in complex biological mixtures. *Journal of the American Statistical Association* **107** 1259–1271.
- BARRETT, P. J., CHEN, J., CHO, M.-K., KIM, J.-H., LU, Z., MATHEW, S., PENG, D., SONG, Y., VAN HORN, W. D., ZHUANG, T., SÖNNICHSEN, F. D. and SANDERS, C. R. (2013). The quiet renaissance of protein nuclear magnetic resonance. *Biochemistry* **52** 1303–20.
- BLOCH, F., HANSON, W. and PACKARD, M. (1946). Nuclear induction. *Phys Rev* **69** 127.
- BRETTTHORST, G. L. (1990). Bayesian analysis. I. Parameter estimation using quadrature NMR models. *Journal of Magnetic Resonance (1969)* **88** 533–551.
- CARDOSO, M. F., SALCEDO, R. and FEYO DE AZEVEDO, S. (1996). The simplex-simulated annealing approach to continuous non-linear optimization. *Computers & chemical engineering* **20** 1065–1080.
- CHEN, Q., MARBLE, A. E., COLPITTS, B. G. and BALCOM, B. J. (2005). The internal magnetic field distribution, and single exponential magnetic resonance free induction decay, in rocks. *Journal of Magnetic Resonance* **175** 300–308.
- DALITZ, F., CUDAJ, M., MAIWALD, M. and GUTHAUSEN, G. (2012). Process and reaction monitoring by low-field NMR spectroscopy. *Progress in nuclear magnetic resonance spectroscopy* **60** 52–70.
- DOU, L. and HODGSON, R. (1995). Bayesian inference and Gibbs sampling in spectral analysis and parameter estimation. I. *Inverse problems* **11** 1069.
- ERNST, R. R. (1992). Nuclear magnetic resonance Fourier transform spectroscopy. *Bioscience reports* **12** 143–187.
- EVILIA, R. F., EFFIONG, R. and WHITTENBURG, S. L. (1993). Bayesian estimation of NMR spectral parameters under low signal-to-noise conditions. *Spectroscopy letters* **26** 1559–1570.
- GEMAN, S. and GEMAN, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6** 721–741.
- GLADDEN, L. (1994). Nuclear magnetic resonance in chemical engineering: principles and applications. *Chemical engineering science* **49** 3339–3408.
- HAACKER, E. M., BROWN, R. W., THOMPSON, M. R. and VENKATESAN, R. (1999). *Magnetic resonance imaging: physical principles and sequence design* **1**. Wiley-Liss New York.
- HUTTON, W. C., BRETTTHORST, G. L., GARBOW, J. R. and ACKERMAN, J. J. (2009). High dynamic-range magnetic resonance spectroscopy (MRS) time-domain signal analysis. *Magnetic Resonance in Medicine* **62** 1026–1035.
- KEELER, J. (2011). *Understanding NMR spectroscopy*. Wiley.
- MALZ, F. and JANCKE, H. (2005). Validation of quantitative NMR. *Journal of pharmaceutical and biomedical analysis* **38** 813–23.

- NELDER, J. A. and MEAD, R. (1965). A simplex method for function minimization. *The computer journal* **7** 308–313.
- PURCELL, E., TORREY, H. and POUND, R. (1946). Resonance absorption by nuclear magnetic moments in a solid. *Phys Rev* **69** 37–38.
- RABI, I., MILLMAN, S., KUSCH, P. and ZACHARIAS, J. (1939). The Molecular Beam Resonance Method for Measuring Nuclear Magnetic Moments. The Magnetic Moments of ${}^6\text{Li}$, ${}^7\text{Li}$ and ${}^{19}\text{F}$. *Physical Review* **55** 526.
- RASMUSSEN, C. E. and GHARAMANI, Z. (2001). Occam’s Razor. In *Advances in Neural Information Processing Systems*.
- RUBTSOV, D. V. and GRIFFIN, J. L. (2007). Time-domain Bayesian detection and estimation of noisy damped sinusoidal signals applied to NMR spectroscopy. *Journal of Magnetic Resonance* **188** 367–379.
- RUBTSOV, D. V., WATERMAN, C., CURRIE, R. A., WATERFIELD, C., SALAZAR, J. D., WRIGHT, J. and GRIFFIN, J. L. (2010). Application of a bayesian deconvolution approach for high-resolution 1h nmr spectra to assessing the metabolic effects of acute phenobarbital exposure in liver tissue. *Analytical chemistry* **82** 4479–4485.
- YOON, J. W., GODSILL, S., KUPČE, E. and FREEMAN, R. (2006). Deterministic and statistical methods for reconstructing multidimensional NMR spectra. *Magnetic Resonance in Chemistry* **44** 197–209.

ANDREW GORDON WILSON
DEPT. OF ENGINEERING
UNIVERSITY OF CAMBRIDGE
CAMBRIDGE, UNITED KINGDOM
E-MAIL: agw38@cam.ac.uk

DANIEL J. HOLLAND
DEPT. OF CHEMICAL ENGINEERING AND BIOTECHNOLOGY
UNIVERSITY OF CAMBRIDGE
CAMBRIDGE, UNITED KINGDOM
E-MAIL: djh79@cam.ac.uk

MICK D. MANTLE
DEPT. OF CHEMICAL ENGINEERING AND BIOTECHNOLOGY
UNIVERSITY OF CAMBRIDGE
CAMBRIDGE, UNITED KINGDOM
E-MAIL: mdm20@cam.ac.uk

YUTING WU
DEPT. OF CHEMICAL ENGINEERING AND BIOTECHNOLOGY
UNIVERSITY OF CAMBRIDGE
CAMBRIDGE, UNITED KINGDOM
E-MAIL: yw316@cam.ac.uk

SEBASTIAN NOWOZIN
MICROSOFT RESEARCH
21 STATION ROAD
CB12FB CAMBRIDGE, UNITED KINGDOM
E-MAIL: Sebastian.Nowozin@microsoft.com

LYNN F. GLADDEN
DEPT. OF CHEMICAL ENGINEERING AND BIOTECHNOLOGY
UNIVERSITY OF CAMBRIDGE
CAMBRIDGE, UNITED KINGDOM
E-MAIL: lfg1@cam.ac.uk

ANDREW BLAKE
MICROSOFT RESEARCH
21 STATION ROAD
CB12FB CAMBRIDGE, UNITED KINGDOM
E-MAIL: Andrew.Blake@microsoft.com