Published in final edited form as:

J Magn Reson. 2006 March ; 179(1): 152–163. doi:10.1016/j.jmr.2005.11.004.

Quantifying in vivo MR spectra with circles *

Refaat E. Gabr^b, Ronald Ouwerkerk^a, and Paul A. Bottomley^{a,b,*}

^a Division of MR Research, Department of Radiology, Johns Hopkins University, Baltimore, MD, USA

^b Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD, USA

Abstract

Accurate and robust quantification of in vivo magnetic resonance spectroscopy (MRS) data is essential to its application in research and medicine. The performance of existing analysis methods is problematic for in vivo studies where low signal-to-noise ratio, overlapping peaks and intense artefacts are endemic. Here, a new frequency-domain technique for MRS data analysis is introduced wherein the circular trajectories which result when spectral peaks are projected onto the complex plane, are fitted with active circle models. The use of active contour strategies naturally allows incorporation of prior knowledge as constraint energy terms. The problem of phasing spectra is eliminated, and baseline artefacts are dealt with using active contours-snakes. The stability and accuracy of the new technique, *CFIT*, is compared with a standard time-domain fitting tool, using simulated ³¹P data with varying amounts of noise and 98 real human chest and heart ³¹P MRS data sets. The real data were also analyzed by our standard frequency-domain absorption-mode technique. On the real data, *CFIT* demonstrated the least fitting failures of all methods and an accuracy similar to the latter method, with both these techniques outperforming the time-domain approach. Contrasting results from simulations argue that performance relative to Cramer-Rao Bounds may not be a suitable indicator of fitting performance with typical in vivo data such as these. We conclude that *CFIT* is a stable, accurate alternative to the best existing methods of fitting in vivo data.

Keywords

Magnetic resonance spectroscopy; Frequency domain quantification; Active circle models; Prior knowledge; Spectral analysis

1. Introduction

Accurate quantification of in vivo MR spectroscopic (MRS) data is often thwarted by low signal-to-noise ratio (SNR), overlapping resonances, and baseline artefacts arising from physiological motion, localizing gradients, and/or delayed acquisition. Standard quantification approaches apply numerical methods to estimate parameters that characterize a mathematical model of the MRS signal, and are generally classified as either time- or frequency-domain methods. Time-domain methods operate directly on the raw data and offer some convenient ways of addressing incomplete or corrupted data [1]. They can be further divided into interactive and black box methods. Interactive methods employ numerical optimization techniques to iteratively solve a non-linear least-squares optimization problem [1–3] based on

*NIH Grant 2RO1 HL 56882.

* Corresponding author. Fax: +1 410 614 1977. E-mail address: bottoml@mri.jhu.edu (P.A. Bottomley).

user-defined initial parameter values. Prior knowledge about peak structure may be used to constrain the solution. The optimization can be performed either locally, or globally [4,5]. Global optimization does not require user interaction but is inherently computationally inefficient. Black box methods provide non-iterative fitting solutions employing linear prediction or state-space theory [1,3,6–8]. These are fast and require little user interaction, but perform poorly with in vivo data because they allow little or no prior knowledge. A major problem with time-domain techniques is that frequency-selective analysis is not straightforward, albeit soluble [9–12]. The frequency domain is certainly the natural domain for spectral analysis: it is in the frequency domain wherein user interactivity invariably occurs, from the input of initial values and prior knowledge, to the final judgment of acceptance of a fit.

Frequency-domain methods model the spectrum as a set of Lorentzian, Gaussian or Voigt peaks, and fit the data using least-squares methods. The real and imaginary parts or the complex signal [13,14] can be used in the fitting process, but quantification is more robust when the number of parameters being estimated is minimized. This can be achieved, for example, by first correcting the spectrum for zero and first-order phase, so that all of the peaks have zero phase. Then, only the pure absorption-mode signal need be fitted. However, the baseline is often severely distorted by the acquisition delays required for spatial encoding gradients which convolve the spectrum with sinc functions [15], by transient eddy currents from the gradients, and by physiological motion. This too must be corrected, either by modeling the baseline with a set of polynomial or sinusoidal basis functions, or by deconvolution.

While new solutions are being investigated [16–18], phase and baseline corrections are difficult to automate and usually cannot match the performance of an experienced manual operator when artefacts are large. This is evidenced, for example, by a 50% increase in the standard deviation of normal human cardiac phosphocreatine (PCr) to adenosine triphosphate (ATP) ratios quantified in phosphorus (^{31}P) spectra by automated Marquardt frequency-domain fitting and by non-linear least squares fitting in the time-domain [19], as compared to manual frequency-domain fitting of data acquired with the identical protocol and equipment [20]. A 50% increase in the uncertainty of a PCr/ATP measurement is equivalent to a 50% reduction in the SNR of the measurement from what could have been achieved. Frequency-domain techniques have been reviewed elsewhere [13], and other variants include the LCMODEL which fits the spectrum with a linear combination of in vitro spectra acquired from individual metabolites [21,22], and time-frequency analysis using wavelets which separates and subtracts each component from the raw signal [23].

In this paper, we present a new frequency-domain technique for MRS quantification that overcomes phasing problems and addresses baseline issues with current techniques as applied to low-SNR, high-artifact regimes. Using a Lorentzian model, we show that the real-imaginary trajectory of each peak in a spectrum scribes a circle in the complex plane with a diameter equal to the peak height. The complex spectrum is fitted using a set of active circle models. The use of both real and imaginary components as a continuous circular trajectory with this circle fitting (CFIT) approach improves the stability of the estimation process, effectively eliminates zero and first-order phase-correction, and reduces the baseline problem to one of completing circles for peaks whose loops are open in the complex plane. Prior knowledge is easily incorporated as external energy terms [24]. The CFIT technique is tested with simulated brain ^{31}P spectra, as well as with real in vivo localized human chest and heart ^{31}P spectra in which problems of low SNR, phasing, baseline, and overlapping peaks are endemic. The stability and accuracy of CFIT are compared with our current standard analysis tool, a frequency-domain absorption mode fitting routine, CSX [25,26], and with a widely available time-domain method, AMARES [6].

2. Theory

The complex time-domain MRS signal, $s(t)$, is most commonly modeled as the sum of exponentially damped sinusoids:

$$s(t) = \sum_{k=1}^K a_k e^{j\varphi_k} e^{(-d_k - j2\pi f_k)t} + n(t), \quad (1)$$

where K is the number of sinusoidal components in the signal, f_k is the frequency of the k th sinusoid, a_k is its amplitude, d_k is its damping constant ($=1/T_{2k}^*$), φ_k is its phase, and $j = \sqrt{-1}$. The term $n(t)$ is a random complex white noise with a circular Gaussian distribution. In the frequency-domain, the spectrum, $S(f)$, is the Fourier transform (FT) of Eq. (1), and is the sum of a series of Lorentzian functions, one for each peak:

$$S(f) = \sum_{k=1}^K \frac{a_k}{d_k} e^{j\varphi_k} \left(\frac{1}{1 + (2\pi(f - f_k)/d_k)^2} + j \frac{2\pi(f - f_k)/d_k}{1 + (2\pi(f - f_k)/d_k)^2} \right) + N(f). \quad (2)$$

Eq. (2) describes a helix in the three-dimensional (3D) space comprised of the real-imaginary complex plane and the chemical shift/frequency axis. The bracketed terms are of the form, $x = 1/(1 + t^2)$, $y = t/(1 + t^2)$, which is a parametric equation for a circle. The projection of the helix onto the complex plane thus forms a set of parametric equations that, for Lorentzian lines, describe circles with centers at the real-imaginary coordinates $(a_k/2d_k, 0)$ and radii $a_k/2d_k$ as exemplified for one peak in Fig. 1A for $\varphi = 0$. The effect of setting $\varphi \neq 0$, is simply to rotate the circle about the origin of the complex plane, moving the circle's center (Fig. 1C). The relationships between the damped sinusoid, the corresponding peak in the spectrum and the circle in the complex plane are summarized in Table 1. The model spectrum is thus a set of circles, one for each peak, which form a series of nested curves in the complex plane resolved in the orthogonal dimension by the f parameter. This is illustrated in Fig. 2, which shows a simulated ^{31}P spectrum from the human brain. Fig. 2A is the real part of the spectrum. Fig. 2B shows the projection of the spectrum onto the complex plane revealing the circular patterns, while Fig. 2C is a 3D plot of the spectrum showing its helical trajectory in this space.

Because the trajectory of each peak in the complex plane is circular, we can model them with active circles. The active circles are contours that adaptively deform to best fit the model spectrum to the measured spectrum by minimizing a measure of the fit error while preserving their circular shape in the complex plane. The fit error, $e(f)$, is the complex difference between the measured spectrum and the sum of the model defined in Eq. (2) plus the baseline in the complex plane. Initial values for the parameters are either provided by the user, estimated from the spectrum, or set to fixed values. The peak center frequency, f_k , needs to be set reasonably accurately, especially when the peaks partially overlap as is typically the case with ^{31}P MRS ATP doublets and triplets. Then, minimizing the *energy*, of the fit error

$$E = \int e(f)^* e(f) df \quad (3)$$

using a gradient descent approach leads to an iterative solution in which the parameters change under the effect of the fit error. In this sense, $e(f)$ acts as an external error *force* [24].

The error force acting on the circumference of an active circle can be resolved into a radial-like component emanating from the origin of the complex plane, $e_{rk}(f)$, and a component orthogonal to it, $e_{ok}(f)$. The component tangential to the circle is denoted $e_{tk}(f)$. The radial-like force exerts a compression/inflation effect that affects the circle radius and hence peak height. The orthogonal force rotates the circle about the complex origin, affecting the circle orientation or peak phase. The tangential force has two effects: a stretching/shrinking effect that changes the point density on the circle and adjusts the linewidth; and rotation of the whole circle about its center, affecting the peak center frequency. These four effects are used to adapt each circle by changing the appropriate parameter so as to minimize the acting force. The error forces along the circle circumference are weighted using a generally bell-shaped function, w , that is centered around the peak center frequency, to emphasize the contribution of nearby error forces in the circle evolution process, reduce the effect of noise and artefacts, and provide a means for selective analysis.

We now describe how the model circles evolve under the effect of the error forces. Details of how the minimization of E leads to the following results are given in Appendix A. Consider the k th circle, the incremental changes in the circle radius r_k , the center frequency f_k , the phase φ_k , and the damping constant d_k are calculated according to the direction of the gradient of E as follows:

$$\Delta r_k = c_r \int w_{rk}(f - f_k) e_{rk}(f) df, \quad (4)$$

$$\Delta f_k = c_f \int w_{fk}(f - f_k) e_{tk}(f) df, \quad (5)$$

$$\Delta \varphi_k = c_\varphi \int w_{\varphi k}(f - f_k) e_{ok}(f) df, \quad (6)$$

$$\Delta d_k = c_d \int w_{dk}(f - f_k) e_{tk}(f) \text{sgn}(f - f_k) df, \quad (7)$$

where $e_{rk}(f)$, $e_{ok}(f)$, and $e_{tk}(f)$ are as defined above, $w_{pk}(\cdot)$ is the weighting function for parameter p in the k th peak, c_r , c_f , c_φ , and c_d are the coefficients determining the step size for each parameter, and sgn is the signum function. The weighting function as well as the coefficients are generally functions of the k th peak being processed, and time. All circles deform during each iteration until they either reach equilibrium, or the number of iterations exceeds a designated threshold. The criterion used for equilibrium is that there is no significant change in the area of all the peaks.

Baseline distortions manifest as a displacement of points on the circle such that the ends of each peak no longer complete a circle in the complex plane. To account for baseline distortion, the baseline is modeled as a free-form active contour: a snake [24]. The snake does not take the form of any specific function such as a polynomial or the like (Appendix A). While constrained by its elastic properties, the snake is adjusted to minimize an energy term representing the residuals of the fit between the measured data and the model spectrum. The elastic properties of the contour are classically defined in terms of its length and curvature. The baseline is simultaneously fitted with the peaks by updating both the active circles and the baseline in each iteration of the fitting process.

The CFIT technique can thus be summarized by the following steps:

1. Initialize K active circles by peak picking or by a black-box method;
2. Calculate the error force between the measured data and the model spectrum;
3. Adjust circle parameters to minimize the error energy for all circles using Eqs. (4)–(7);
4. Adjust the baseline snake parameters to minimize the fit error;
5. Repeat steps 2–4 until either equilibrium or the maximum number of iterations is reached;
6. Output peak frequencies, areas, and heights. Display the fit.

3. Methods

3.1. Preprocessing

CFIT is applied to data prepared in the frequency-domain in the usual manner with exponential filtering, zero-filling (ZF), and FT of the time-domain signal [27]. The spectrum is typically normalized by the maximum value of the magnitude of a standard spectrum so that the same set of parameters, c_r , c_f , c_ϕ , and c_d , can be used for all spectra acquired under similar conditions. Note that in Fig. 1, the portion of the peak with the highest SNR includes most of the circle while the relatively long peak tails with low SNR contribute to a very small fraction of the circle's arc. While the number of points added by zero-padding is uniform across the spectrum, the weighting function, w , used in the fitting process, emphasizes the contribution of points with high SNR. Accordingly, simulations showed that increasing the ZF factor increases the number of sample points in the high-SNR portion of the circle arc, and improves both the fitting success rate, and reduces the relative root mean square error (RRMSE) and the root mean square error (RMSE) of the amplitude, damping and frequency estimates, as shown in Fig. 3.

Although phasing the spectrum is not required in CFIT, a basic zero- and first-order phasing function is included for the sole purpose of facilitating a user's ability to interactively pick peak parameters. This is standard for most current user-interactive approaches. The fitting process actually uses the initial un-phased spectrum to avoid any added distortions that may arise from first-order phase correction [15]. Phasing may also be applied to the post-processed fitted spectrum, if desired, for display purposes.

3.2. Adding prior knowledge

Prior knowledge is easily incorporated in CFIT by assigning an energy term to each constraint or feature that one wishes the solution to stick to. The constraints need not be linear. Minimization of the constraint energy is also carried out using the energy descent method. The external constraints affecting the model spectrum can be represented by spring forces that drive the contour towards the features of interest. Each of the active circles adapts under the influence of the error forces and the constraints until it is stabilized. In general, a parameter p that is to be estimated, is adjusted in each iteration under the influence of the two sets of forces via

$$p \leftarrow p + k_p(p_d - p) + \int e_p(f) df, \quad (8)$$

where p_d is the desired value that p should approach, e_p is the external error forces and k_p is the spring constant. If the constraint to be imposed is a ratio m between two parameters p and q , such as the peak amplitudes in a multiplet, then we set $p_d = mq$. The q parameter is then also

affected by the prior knowledge by setting $q_d = p/m$. If a fixed distance Δp between two parameters p and q is required, like the fixed frequency separation between the ATP fine peaks, we set $p_d = q + \Delta p$ and $q_d = p - \Delta p$. Other constraints can be imposed in a similar fashion.

4. Results

4.1. Fitting simulated MRS signals

CFIT was first tested with a simulated in vivo ^{31}P human brain spectrum consisting of 256 complex time-domain data points, zero-filled eight-times and comprised of 11 exponentials, defined elsewhere [6] with parameters summarized in Table 2. The spectrum in Fig. 2 is illustrated with complex white Gaussian noise of standard deviation (SD) $\sigma = 25$ added, and a 10-Hz exponential filter applied. The sampling frequency, f_s , is 3 kHz, and the fine peaks in the ATP multiplets are split by 16 Hz. The technique was tested on 100 different simulated signal-plus-noise realizations, for each of 20 SNR levels in the range $5 \leq \sigma \leq 100$ with a step of 5, for a total of 2000 runs. The initial values for the frequency and damping ratio were picked once at each noise level and the same initial values were used for all signals at the same noise level. To obtain a good initial value for the circle phase, a small number of points around the chosen center frequency were fitted to a circle by a standard least-squares fitting routine. The phase of the fitted circle was taken as the initial phase for each circle. Although an accurate initial value is not required for phase, care must be taken not to start with an initial value for the circle/helix that is 180° from the true phase.

Prior knowledge of the ATP doublets (α -ATP and γ -ATP peaks) and triplet (β -ATP) was implemented via the following constraints; $d_1 = d_2 = d_3$, $d_4 = d_5$, $d_6 = d_7$, $a_1 = a_2/2 = a_3$, $a_4 = a_5$, $a_6 = a_7$, $\phi_1 = \phi_2 = \phi_3$, $\phi_4 = \phi_5$, $\phi_6 = \phi_7$, $f_3 = f_2 + 16 = f_1 + 32$, $f_5 = f_4 + 16$, and $f_7 = f_6 + 16$. The parameters c_r , c_f , c_ϕ , and c_d were empirically set to 0.65, 0.4, 1.6, and 41, and k_r , k_f , k_ϕ , and k_d to 0.2, 0.8, 0.4, and 0.6, respectively, to optimize the fit. The weighting functions are chosen as Lorentzian that evolve during peak processing, as described in Appendix A. The Lorentzian window is chosen to have a linewidth that is half the line width of the currently processed peak at the current iteration.

For comparison, the same data was quantified using AMARES as provided in the standard jMRUI package, available from <http://www.mrui.uab.es/mrui/upon request>. AMARES [6] is a time-domain fitting technique that has been demonstrated to have stable and accurate performance [6,8]. This program requires the user to pick the peaks providing initial values for the peak frequencies and linewidths, and to set the constraints in a similar fashion as in CFIT. The constraints used for AMARES (d_k , a_k , f_k , and ϕ_k) and CFIT were identical.

As reported in [8], the robustness of each method is evaluated by computing its success rate, i.e., the number of times, out of the total number of simulation runs, that the method is able to resolve the 11 peaks within specific intervals lying symmetrically around the true frequencies of the peaks. The half-widths of these intervals are set to 8 Hz, i.e., half the separation of the closest peaks in the data. The peak error is quantified by the mean RRMSE defined as

$$\text{mean RRMSE} = \frac{100}{K} \sum_{k=1}^K \sqrt{\frac{1}{J} \sum_{j=1}^J \frac{(p_{kj} - p_k)^2}{p_k^2}} \quad (9)$$

for the amplitude and damping estimates, and as the mean RMSE

$$\text{mean RMSE} = \frac{1}{K} \sum_{k=1}^K \sqrt{\frac{1}{J} \sum_{j=1}^J (p_{kj} - p_k)^2} \quad (10)$$

for the frequency estimates, where $K = 11$, is the total number of peaks, J is the number of simulation runs in which the method was able to find every peak within the corresponding frequency interval, p_{kj} denotes the parameter estimate of the k th peak obtained in simulation run j , and p_k denotes the true parameter value of the k th peak [8].

Fig. 3 shows that the success rate for CFIT with high zero-filling and AMARES using simulated data are comparable. Both perform perfectly when SNR is high ($\sigma < 50$), but deteriorate as SNR declines further. Figs. 3(B–D) plots the mean RRMSE values for the amplitude and the damping constant, and the mean RMSE values for the frequency of the ATP complex (the first seven peaks in the simulated spectrum) as a function of noise level for the simulated data. The Cramer-Rao bound (CRB), a measure of the best attainable precision of an unbiased estimate in the presence of random noise [28], is also plotted for each of the estimated parameters. The results of these simulations are consistent with AMARES' reputation as one of the best performing methods relative to the CRB [8]. Even so, relative performance can reverse in the presence of baseline artifacts. For example, perturbing the first four time points of the model brain spectrum by multiplication by random numbers produces a mild baseline roll as shown in Fig. 4. The success rate of CFIT is comparable to the rate without perturbation, albeit with some increase in the amplitude error. However the performance for uncompensated AMARES is decimated. AMARES performance does improve with initial filtering to reverse the effect of the imposed distortion, but CFIT remains the more successful in this instance (Fig. 4A). Fig. 5 shows a sample output of CFIT for one of the processed simulated spectrum for the case $\sigma = 25$. At this level of noise, the SNR of the β -ATP peak is about 25.

4.2. In vivo ^{31}P data

CFIT and AMARES were tested on 98 in vivo localized ^{31}P spectra acquired in 5–7 min from the human chest and heart at 1.5 T using a cardiac-gated, quantitative, one-dimensional chemical shift imaging (CSI) protocol employing adiabatic 60° excitations and an 0.7 ms acquisition delay as described elsewhere [29]. The spectra derive from normal subjects (30%), patients with dilated cardiomyopathy and congestive heart failure (40%), and patients with left ventricular hypertrophy (30%), participating in ongoing studies of cardiac creatine kinase flux wherein scan times are limited by the need to complete six MRS acquisitions in a single exam [29,30]. Despite its performance in simulations with respect to the CRB, AMARES is not used for in vivo cardiac ^{31}P MRS quantification in our laboratory due to an unacceptable failure rate. Instead, we routinely use a frequency-domain interactive fitting program, CSX [25,26], wherein peak height, frequency, and linewidths are fitted by a simplex routine after the user phases the spectrum, selects points for baseline correction, and identifies initial frequencies of peaks selected for fitting. The performance of CFIT was therefore additionally compared to CSX, our reference standard, for which fit results for all 98 spectra had already been documented. Spectra were thus selected for analysis based only on the chronological order in which studies were performed, and the availability of pre-existing CSX measurements with a minimum fitting threshold for SNR of about 5 for β -ATP.

Spectra to be analyzed for CFIT were prepared with 8-fold zero-filling, and a 10-Hz exponential filter. The adapting parameters and the windowing function were the same as in the simulations. The spectra were roughly phased by an experienced in vivo spectroscopist (PAB) who identified the initial peak locations. As in the simulations, the β -ATP peak was fitted to a triplet and the α - and γ -ATP peaks were fitted to doublets. One-to-three Lorentzian peaks were used

to model the PCr peak since its shape was not Lorentzian: the use of multiple Lorentzian peaks to approximate other non-Lorentzian peak shapes (Gaussian, Voigt, etc) has been reported before [7]. The multiple PCr peaks are linked so that they have the same phase. This constraint is essential for eliminating solutions which have peaks that cancel. The in vivo spectra were quantified with AMARES in the same way by the same user. AMARES requires additional prior knowledge estimates of linewidths for the ATP multiplets, which are provided interactively. The spectra were analyzed 0.5–3 years earlier in the same way by the same user with the CSX program, except that they were more carefully phased (zero and first-order), baselines were corrected with user-defined points between the peaks, and the ATP peaks were fitted only as singlets but with Gaussian-shaped lines. These differences reduced the number of parameters that needed to be estimated for the CSX method compared to CFIT and AMARES.

Because the true values of in vivo parameters are not known, fitting failures cannot in practice be defined by comparison with true values. However, we can record a failure when the fit misses one peak completely, or when anomalously small or anomalously large areas of a peak are returned. The number of fit failures involving one or more of the ATP or PCr peaks in the ^{31}P MRS cardiac data was 4/98 for CFIT. This compares with 29/98 for AMARES, and 8/98 failures for CSX, the latter involving ATP only. The total number of cases in which at least one of the three fitting methods fails was 31.

The ratio of the peak areas of the β -ATP and γ -ATP provides an independent gauge of the robustness of the different fitting routines. This ratio must be approximately unity since the nucleoside triphosphate sources are the same, and additional contributions to the γ -resonance from other phosphates are small. The ratio obtained using the three fitting methods is plotted in Fig. 6, and the means and SD for the ratio obtained by each method are listed in Table 3. While AMARES yields erroneous results that greatly deviate from unity in many cases, CFIT performs comparably with CSX, our current standard.

The PCr/ATP ratio of peak areas is a physiologically important metabolic measure [20]. Fig. 7 plots the PCr/ β -ATP ratio measured by AMARES and CFIT against our current standard, CSX. With perfect agreement, all the data would fall on the 45° line and lie in the known range of about 0.2–5 for normal and sick human chest and heart. The CSX data show reasonable agreement with CFIT but much larger disparity for AMARES, with values lying outside the physiological range. Table 4 reports the Pearson correlation coefficient analysis of the PCr/ β -ATP, PCr/ α -ATP, and PCr/ γ -ATP ratios measured using AMARES, CSX, and CFIT. To avoid degrading the regression with the much larger number of outliers produced by AMARES, PCr/ATP ratios greater than 10 were eliminated from this analysis. Clearly, the agreement between CSX and CFIT is the highest for all three ratios. Again even with all of the AMARES fit failures excluded, assuming CSX is our standard, the mean RRMSE of the peak areas of β -ATP, α -ATP, and PCr peaks show greater scatter when measured with AMARES than with CFIT, especially for PCr, as demonstrated in Table 5. A sample CFIT analysis of one of the processed spectra is shown in Fig. 8.

The computational time in CFIT is longer than AMARES but comparable or faster than the simple CSX frequency-domain fitting routine. Implementation on a PC (Intel Pentium M processor 1.5 GHz and 512 MB of memory) operating on the Microsoft Windows XP platform in MATLAB v6.5 (Mathworks, Natick, MA) required an average of 87 ± 47 iterations to achieve a stable solution in 2.9 ± 1.6 s.

5. Discussion

We have introduced a new frequency-domain spectral analysis technique, CFIT, that utilizes *active circles* to directly fit the circular trajectory that Lorentzian resonances scribe in the complex plane. The method has the advantage of eliminating the effect of zero- and first-order phase correction on peak quantification. Baseline roll and overlapping peaks affect the continuity of circles in the complex plane, and are dealt with using *active snakes*. Prior knowledge is incorporated in the usual way, and non-Lorentzian peaks accommodated using multiple Lorentzian peaks. The ability of CFIT to accommodate large first-order phase errors and baseline artefacts suits it well to in vivo analyses where such problems are endemic. It is especially suited to quantifying low-SNR high-artifact free-induction decays in which acquisition is delayed by gradient-encoding, such as cardiac ^{31}P CSI data [29,30].

It is disappointing that the performance of CFIT is inferior to that of the standard time-domain fitting technique, AMARES, with simulated data, as indexed by the Cramer-Rao bounds for key fitting parameters in the absence of baseline artefacts. However, AMARES is not used in our laboratory for analyzing human cardiac ^{31}P CSI data due to its unacceptably high failure-to-fit rate, an anecdotal observation not previously documented. Therefore, the performance comparison of CFIT on about a hundred real in vivo data sets analyzed by our standard absorption-mode frequency-domain technique, CSX, as well as by AMARES is, in our view, the more relevant comparison. Despite the excellent performance of AMARES on simulated noisy data free of baseline artefacts, we document with simulations that the failure rate of AMARES can become excessive in the presence of distortions that cause baseline artifacts. Moreover, for the in vivo data its failure rate is indeed unacceptably high compared with the other two methods, and its accuracy much lower. Conceivably, better adjustment of AMARES constraints and filtering could enhance the performance with our in vivo data, just as better filtering that does not distort the circles could also benefit CFIT. In our view, the technique of choice is the one that provides the best fit to the worst spectrum from which measurements *must* be made, within the limits of detection permitted by the SNR. CFIT had the lowest failure rate with real human cardiac ^{31}P CSI data, and performed comparable to CSX with respect to measures of accuracy ($\beta\text{-ATP}/\gamma\text{-ATP}$; Fig. 6) and scatter (RRMSE; Tables 3–5, Figs. 6 and 7) on peaks with SNR as low as ~ 5 . By these criteria CFIT is a suitable replacement for CSX, our current standard for cardiac ^{31}P MRS.

New quantification techniques are often accompanied by claims of accuracy and robustness [6,14]. However, caution is warranted when advantages are demonstrated only in cases of high SNR, idealized line-shapes, or signals containing only Gaussian noise, such as is common with simulations. In particular, many in vivo signals contain significant artefacts that cannot be modeled with Gaussian noise. An estimator that deals with in vivo data assuming Gaussian noise may thus be misled in situations where such artefacts arise. This observation together with our performance measurements on simulated and real data, render questionable the suitability of the Cramer-Rao bound as a measure of the performance of estimators of in vivo MRS signals in such settings.

Two key factors help CFIT overcome the effects of artefacts, resulting in better in vivo performance. The first is that the determination of step size during parameter adaptation involves an integration process, which tends to cancel random errors. The second factor is the bell-shaped window (implemented here as Lorentzian) used for weighting the error force that affects each active circle. Such weighting greatly reduces the contamination of nearby artefacts and also allows for selective frequency analysis. In this case, residuals from outside the quantified area are detected using the baseline snake. The ability to impose prior knowledge is an additional major factor that stabilizes the performance of fitting routines in general, including CFIT. Incorporating prior knowledge has been shown to benefit accuracy [28], and

a wide range of prior knowledge including non-linear constraints, can be incorporated to guide quantification with CFIT as desired.

In conclusion, CFIT is a new frequency-domain spectral fitting technique that has been tested on real in vivo cardiac ^{31}P MRS data, warts-and-all. It provides accurate stable performance, and appears to be a suitable alternative to existing standards.

Acknowledgments

We thank Roger Eastman of Loyola College in Baltimore and Teresa Laudadio, Department of Electrical Engineering, ESAT-SCD(SISTA), Katholieke Universiteit Leuven, Leuven (Belgium) for help and discussions. The data derive from ongoing human heart studies performed with Robert G Weiss, for which Peter Barker's CSX analysis tools have served us well and long: both are at Johns Hopkins.

References

1. Vanhamme L, Sundin T, Hecke PV, Huffel SV. MR spectroscopy quantitation: a review of time-domain methods. *NMR Biomed* 2001;14(4):233–246. [PubMed: 11410941]
2. van der Veen JWC, de Beer R, Luyten PR, van Ormondt D. Accurate quantification of in vivo ^{31}P NMR signals using the variable projection method and prior knowledge. *Magn Reson Med* 1988;6(1):92–98. [PubMed: 3352510]
3. Vanhamme L, Van Huffel S, Van Hecke P, van Ormondt D. Time domain quantification of series of biomedical magnetic resonance spectroscopy signals. *J Magn Reson* 1999;140(1):120–130. [PubMed: 10479554]
4. DiGennaro FS, Cowburn D. Parametric estimation of time-domain NMR signals using simulated annealing. *J Magn Reson* 1992;96:582–588.
5. Metzger GJ, Patel M, Hu X. Application of genetic algorithms to spectral quantification. *J Magn Reson B* 1996;110(3):316–320. [PubMed: 8867450]
6. Vanhamme L, van den Boogaart A, Van Huffel S. Improved method for accurate and efficient quantification of MRS data with use of prior knowledge. *J Magn Reson* 1997;129(1):35–43. [PubMed: 9405214]
7. van den Boogaart A. Quantitative data analysis of in vivo MRS data sets. *Magn Reson Chem* 1997;35:S146–S152.
8. Laudadio T, Selen Y, Vanhamme L, Stoica P, Van Hecke P, Van Huffel S. Subspace-based MRS data quantitation of multiplets using prior knowledge. *J Magn Reson* 2004;168(1):53–65. [PubMed: 15082249]
9. Sandgren N, Selén Y, Stoica P, Li J. Parametric methods for frequency-selective MR spectroscopy—a review. *J Magn Reson* 2004;168(2):259–272. [PubMed: 15140436]
10. Knijn A, de Beer R, van Ormondt D. Frequency-selective quantification in the time domain. *J Magn Reson* 1992;97:444–450.
11. Vanhamme L, Sundin T, Van Hecke P, Van Huffel S, Pintelon R. Frequency-selective quantification of biomedical magnetic resonance spectroscopy data. *J Magn Reson* 2000;143(1):1–16. [PubMed: 10698640]
12. Romano R, Motta A, Camassa S, Pagano C, Santini MT, Indovina PL. A new time-domain frequency-selective quantification algorithm. *J Magn Reson* 2002;155(2):226–235. [PubMed: 12036333]
13. Mierisova S, Ala-Korpela M. MR spectroscopy quantitation: a review of frequency domain methods. *NMR Biomed* 2001;14(4):247–259. [PubMed: 11410942]
14. Martin YL. A global approach to accurate and automatic quantitative analysis of NMR spectra by complex least-squares curve fitting. *J Magn Reson A* 1994;111:1–10.
15. Allman T, Holland GA, Lenkinski RE, Charles HC. A simple method for processing NMR spectra in which acquisition is delayed: applications to in vivo localized ^{31}P NMR spectra acquired using the DRESS technique. *Magn Reson Med* 1988;7(1):88–94. [PubMed: 3386525]
16. Golotvin S, Williams A. Improved baseline recognition and modeling of FT NMR spectra. *J Magn Reson* 2000;146(1):122–125. [PubMed: 10968964]

17. Soher BJ, Maudsley AA. Evaluation of variable lineshape models and prior information in automated ¹H spectroscopic imaging analysis. *Magn Reson Med* 2004;52(6):1246–1254. [PubMed: 15562473]
18. Stoch G, Olejniczak Z. Missing first points and phase artifact mutually entangled in FT NMR data-noniterative solution. *J Magn Reson* 2005;173(1):140–152. [PubMed: 15705522]
19. Bottomley PA, Atalar E, Weiss RG. Human cardiac high-energy phosphate metabolite concentrations by 1D-resolved NMR spectroscopy. *Magn Reson Med* 1996;35(5):664–670. [PubMed: 8722817]
20. Bottomley PA, Weiss RG, Hardy CJ, Baumgartner WA. Myocardial high-energy phosphate metabolism and allograft rejection in patients with heart transplants. *Radiology* 1991;181:67–75. [PubMed: 1887057]
21. Provencher SW. Estimation of metabolite concentrations from localized in vivo proton NMR spectra. *Magn Reson Med* 1993;30(6):672–679. [PubMed: 8139448]
22. Provencher SW. Automatic quantitation of localized in vivo ¹H spectra with LCModel. *NMR Biomed* 2001;14(4):260–264. [PubMed: 11410943]
23. Serrai H, Senhadji L, de Certaines JD, Coatrieux JL. Time-domain quantification of amplitude, chemical shift, apparent relaxation time T₂, and phase by wavelet-transform analysis. application to biomedical magnetic resonance spectroscopy. *J Magn Reson* 1997;124(1):20–34. [PubMed: 9424306]
24. Kass M, Witkin A, Terzopoulos D. Snakes: active contour models. *Intl J Comput Vision* 1988;2(3):321–331.
25. Soher BJ, van Zijl PCM, Duyn JH, Barker PB. Quantitative proton spectroscopic imaging of the human brain. *Magn Reson Med* 1996;35(3):356–363. [PubMed: 8699947]
26. Bonekamp, D.; Barker, PB. MR spectroscopic imaging software for the neuroradiologist. ASNR, 43rd Annual Meeting; Toronto, CA. 2005.
27. in't Zandt H, van Der Graaf M, Heerschap A. Common processing of in vivo MR spectra. *NMR Biomed* 2001;14(4):224–232. [PubMed: 11410940]
28. Cavassila S, Deval S, Huegen C, van Ormondt D, Graveron-Demilly D. Cramer-Rao bound expressions for parametric estimation of overlapping peaks: influence of prior knowledge. *J Magn Reson* 2000;143(2):311–320. [PubMed: 10729257]
29. Weiss RG, Gerstenblith G, Bottomley PA. ATP flux through creatine kinase in the normal, stressed, and failing human heart. *Proc Natl Acad Sci USA* 2005;102(3):808–813. [PubMed: 15647364]
30. Smith C, Bottomley PA, Gerstenblith G, Schulman SP, Weiss RG. Reduced ATP synthesis through creatine kinase in failing and non-failing hypertrophied human myocardium. *Circulation* 2004;110(III610)
31. Cootes, TF.; Taylor, CJ. Active shape models—‘smart snakes’. *Proc British Machine Vision Conference*; Springer-Verlag. 1992. p. 266-275.
32. Lakshmanan S, Grimmer D. A deformable template approach to detecting straight edges in radar images. *IEEE Trans Pattern Anal Mach Intell* 1996;18(4):438–443.
33. Staib LH, Duncan JS. Boundary finding with parametrically deformable models. *IEEE Trans Pattern Anal Mach Intell* 1992;14(11):1061–1075.

Appendix A

A.1. Active circle models

Consider the model spectrum $S(f)$ given in Eq. (2) including the baseline model. We define:

$$p_k(f) = 2\pi(f - f_k)/d_k \text{ and } r_k = \frac{a_k}{2d_k}. \quad (\text{A.1})$$

Then

$$S(f) = \sum_{k=1}^K 2r_k e^{j\varphi_k} \left(\frac{1}{1+p_k(f)^2} + j \frac{p_k(f)}{1+p_k(f)^2} \right) + \text{baseline}(f) \quad (\text{A.2})$$

The fit error is the difference between the measured and the model spectra,

$$e(f) = S_{\text{measured}}(f) - S(f), \quad (\text{A.3})$$

with the error energy, E , given by Eq. (3). E is minimized using the gradient descent method.

The gradient of E with respect to r_k ,

$$\frac{\partial E}{\partial r_k} \int \frac{\partial}{\partial r_k} e(f) * e(f) df \quad (\text{A.4})$$

reduces to

$$\begin{aligned} \frac{\partial E}{\partial r_k} &= -2 \int \langle e(f), u_k(f)/r_k \rangle df \text{ with } u_k(f) \\ &\equiv 2r_k e^{j\varphi_k} \left(\frac{1}{1+p_k(f)^2} + j \frac{p_k(f)}{1+p_k(f)^2} \right), \text{ or} \end{aligned} \quad (\text{A.5})$$

$$\frac{\partial E}{\partial r_k} = -4 \int \left\langle \frac{1}{1+p_k(f)^2} e(f), \widehat{u}_k(f) \right\rangle df, \quad (\text{A.6})$$

where $\langle \cdot, \cdot \rangle$ is the inner product; $u_k(f)$ and $\widehat{u}_k(f)$ are the radial vector connecting the origin of the complex plane to a point at frequency f on the k th circle and the unit vector in that direction, respectively, and $1/(1+p_k(f)^2)$ serves as the error weighting function, $w_{rk}(f)$, which is maximum at f_k .

The gradient of E with respect to φ_k , f_k , and d_k can be similarly derived as

$$\frac{\partial E}{\partial \varphi_k} = -4 \int \left\langle \frac{r_k}{1+p_k(f)^2} e(f), \widehat{u}_k^\perp(f) \right\rangle df, \quad (\text{A.7})$$

$$\frac{\partial E}{\partial f_k} = -8\pi \int \left\langle \frac{r_k/d_k}{1+p_k(f)^2} e(f), \widehat{\rho}_k^\perp(f) \right\rangle df, \quad (\text{A.8})$$

and

$$\frac{\partial E}{\partial d_k} = -8\pi \int \left\langle \frac{(r_k/d_k)p_k(f)}{1+p_k(f)^2} e(f), \widehat{\rho}_k^\perp(f) \right\rangle df, \quad (\text{A.9})$$

where $u_k^\perp(f) \equiv ju_k(f)$ is perpendicular to the radial vector $u_k(f)$ at f , $\rho_k(f) = 2[u_k(f) - (1/2, 0)e^{j\varphi_k}]$ is the radial vector connecting the center of the k th circle to the point on that circle at frequency f , and $\rho_k^\perp(f) \equiv -j\rho_k(f)$ is thus tangent to the circle at f .

We define the components of $e(f)$ in the radial-like direction $u_k(f)$, the orthogonal direction $u_k^\perp(f)$, and the tangential direction $\rho_k^\perp(f)$ as $e_{rk}(f)$, $e_{ok}(f)$, and $e_{tk}(f)$, respectively. The gradients of E in these directions are

$$\frac{\partial E}{\partial r_k} = -4 \int w_{rk}(f) e_{rk}(f) df, \quad (\text{A.10})$$

$$\frac{\partial E}{\partial \varphi_k} = -4 \int w_{\varphi k}(f) e_{ok}(f) df, \quad (\text{A.11})$$

$$\frac{\partial E}{\partial f_k} = -8\pi \int w_{fk}(f) e_{tk}(f) df, \quad (\text{A.12})$$

and

$$\frac{\partial E}{\partial d_k} = -8\pi \int w_{dk}(f) \text{sgn}(f - f_k) e_{tk}(f) df, \quad (\text{A.13})$$

where the weighting functions $w_{rk}(f)$, $w_{\varphi k}(f)$, $w_{fk}(f)$, and $w_{dk}(f)$ are given by

$$w_{rk}(f) = \frac{1}{1+p_k(f)^2}, \quad (\text{A.14})$$

$$w_{\varphi k}(f) = \frac{r_k}{1+p_k(f)^2}, \quad (\text{A.15})$$

$$w_{fk}(f) = \frac{r_k/d_k}{1+p_k(f)^2}, \quad (\text{A.16})$$

and

$$w_{dk}(f) = \frac{(r_k/d_k) \cdot |p_k(f)|}{1 + p_k(f)^2}. \quad (\text{A.17})$$

These weighting functions are Lorentzian with the same linewidth as the processed peak. In our implementation, we used this same weighting function with half the line width because it yielded better accuracy and stability than with the full linewidth. Other weighting functions with different linewidths and properties can be substituted to achieve preferred levels of artefact rejection.

We refer to this as an *active circle model*. It can be considered an example of the more general set of analytical deformable models that are contours with a general predefined shape. In these models, prior knowledge about the shape is used directly by defining a curve described by a set of parameters that encode the general shape characteristics [31–33].

A.2. Snakes

Active contour models, also known as snakes, are energy-minimizing contours that are constrained by their internal elastic forces, while driven by external forces towards desired external features. In image analysis applications, where active contour models were introduced [24], the external forces are provided by the image brightness and edge characteristics as well as higher-level information. The contour energy functional is constructed so that local minima correspond to desired features. The fitting problem is formulated as an energy minimization problem, and is solved using a suitable minimization technique. A contour is described parametrically by $v(s) = (x(s), y(s))$ where $x(s)$, $y(s)$ are x and y coordinates along the contour and s in $[0,1]$ is the contour length. The contour model defines the energy of a contour $v(s)$, the snake energy E_{snake} , as,

$$E_{\text{snake}} = \int_0^1 E_{\text{int}}(v(s)) + E_{\text{ext}}(v(s)) + E_{\text{con}}(v(s)) ds, \quad (\text{A.18})$$

where E_{int} is the internal energy of the contour, imposing continuity and curvature constraints, E_{ext} is the external energy constructed to attract the snake to desired features; and E_{con} is another energy term which allows various constraints to be applied during the minimization process. An initial contour evolves by minimizing the snake energy using a minimization technique like gradient descent, as used here. In this work, the baseline is modeled as a snake constrained by its smoothness and is driven by the fit error. The baseline is simultaneously fitted with the peaks by performing the active circles updates and a baseline update in each iteration of the solution.

Fig. 1.

The spectrum of a single peak is a helix in the 3D space comprised of the real, the imaginary, and the frequency axes. The projection of the helix onto the complex plane is a circle (A), while the projection of the helix onto the plane of the real and the frequency axes is the absorption-mode spectrum with phase $\varphi = 0^\circ$ (B). Shifting the phase by $\varphi = 60^\circ$ merely rotates the circle diameter about (0,0) by 60° without altering the fundamental properties of the circle (C). However, the spectrum (D) is now a mixture of absorption and dispersion modes and would need to be phased.

Fig. 2.

(A) The real part of the simulated ^{31}P human brain spectrum with noise $\sigma = 25$. Part (B) is the projection of the 3D spectrum trajectory onto the real-imaginary plane. Part (C) is a 3D plot of the spectrum showing the helical trajectory of the peaks.

Fig. 3.

Results of quantifying the simulated ^{31}P spectra comprised of 256 complex points using CFIT with an 8-fold (ZF = 8) and a 4-fold zero-filling (ZF = 4), and AMARES. The success rate (A), the frequency RMSE (B), the amplitude RRMSE (C), and the damping constant RRMSE (D) for the $K = 7$ ATP resonances, are plotted against the SD of the added noise. The Cramer-Rao bound (CRB) on the error for the ATP is also shown in (B–D).

Fig. 4.

Effect of simulated baseline errors on CFIT (ZF = 8) and AMARES performance (A and B) on the 256-point simulated ^{31}P spectra brain spectra from Fig. 3. The signal was distorted by multiplying the first four time points by the random numbers, 0.1453, 0.4715, 1.316, and 0.8948, to produce a mild baseline roll (C). The amplitude error estimates for AMARES in (B) are limited by its low success rate (<10%) in (A). Application of 20-point initial weighting (s-shaped) filter to AMARES (+ filter) significantly improves its success rate (A, center curve).

Fig. 5. Analysis results of the simulated ^{31}P human brain spectrum from Fig. 2 ($\sigma = 25$) using CFIT. From bottom up, the real part of the simulated spectrum, the real part of the fitted spectrum, the real part of the individual circles, and the real part of the residual error (top).

Fig. 6.

The ratio of the peak areas $\beta\text{-ATP}/\gamma\text{-ATP}$ for the in vivo ^{31}P spectra from the human chest and heart, as estimated using CSX, AMARES, and CFIT. The plot excludes all of the 31 fit failures. The true value for the ratio is ~ 1 .

Fig. 7.

The ratio of the peak areas PCr/β-ATP for the in vivo ^{31}P chest and heart spectra as estimated using AMARES (circles) and CFIT (points) vs. the CSX measurements (horizontal axis). AMARES and CFIT points that agree with CSX should fall on the 45° line (dashed). The plot excludes the 31 fit failures. True physiological values for this ratio in human chest and heart lie in the range 0.2–5.

Fig. 8.

Analysis of a real ^{31}P spectrum from a female heart patient using CFIT. The SNR of the β -ATP peak is about 8. From bottom up, the real part of the measured spectrum, the real part of the fitted spectrum, the real part of the individual circles, and the real part of the residual error (top).

Table 1

Relationship between the peak parameters in different domains

Parameter	Time	Frequency	Circle
a_k	Amplitude	Peak area	Radius times number of points in semicircle centered at f_k
d_k	Damping constant	Linewidth	Number of points in semicircle centered at f_k
φ_k	Phase shift	Mode	Circle position/orientation
f_k	Frequency	Peak position	Location of point density minimum on circle
a_k/d_k	Amplitude/decay	Peak height	Circle diameter

Table 2

True parameter values of the simulated ^{31}P MRS signal [6]

k	f_k (Hz)	f_k (ppm)	d_k (Hz)	a_k (a.u.)	ϕ_k ($^\circ$)
1	-86	-17.2	50	75	135
2	-70	-16.6	50	150	135
3	-54	-16.0	50	75	135
4	152	-8.0	50	150	135
5	168	-7.4	50	150	135
6	292	-2.6	50	150	135
7	308	-2.0	50	150	135
8	360	0	25	150	135
9	440	3.1	285.7	1400	135
10	490	5	25	60	135
11	530	6.6	200	500	135

Table 3

Mean and SD of the ratio of the peak area of the β -ATP/ γ -ATP estimated using different methods for the in vivo ^{31}P MRS chest and heart data

Method	Mean	SD
CSX	0.97	0.32
AMARES	1.43	1.09
CFIT	0.84	0.35

The analysis excludes data from the total of 31 fit failures.

Table 4

Correlations (r) between AMARES, CSX, and CFIT for the peak area ratios PCr/ β -ATP, PCr/ α -ATP, and PCr/ γ -ATP for the in vivo ^{31}P MRS signals

Ratio	AMARES-CFIT	AMARES-CSX	CFIT-CSX
PCr/ β -ATP	0.54	0.53	0.75
PCr/ α -ATP	0.36	0.32	0.71
PCr/ γ -ATP	0.60	0.68	0.82

Data from the total of 31 fit failures are eliminated from the regression.

Table 5

Mean RRMSE values and minimum (m) and maximum (M) error relative to CSX for the amplitude estimates obtained using AMARES and CFIT from the in vivo ^{31}P MRS signals after exclusion of the 31 fit failures

Peak	Method	RRMSE [m, M] (%)
β -ATP	AMARES	30.6 [0.29, 81]
	CFIT	25.4 [0.03, 58]
α -ATP	AMARES	69.8 [2.3, 262]
	CFIT	58.9 [2.1, 178]
γ -ATP	AMARES	45.3 [1.7, 93]
	CFIT	60.8 [0.67, 254]
PCr	AMARES	74.5 [2.4, 477]
	CFIT	43.3 [0.04, 93]