Application of Compressed Sensing to Single Voxel J Resolved Magnetic Resonance Spectroscopy: Simulation and In Vitro Results

APPLICATION OF COMPRESSED SENSING TO SINGLE VOXEL J RESOLVED MAGNETIC RESONANCE SPECTROSCOPY: SIMULATION AND IN VITRO RESULTS

BY

BENJAMIN GERAGHTY, B.Eng.

A THESIS

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING

AND THE SCHOOL OF GRADUATE STUDIES

OF MCMASTER UNIVERSITY

IN PARTIAL FULFILMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE

Master of Applied Science (2013) McMaster University

(Electrical & Computer Engineering) Hamilton, Ontario, Canada

TITLE:	Application of Compressed Sensing to Single Voxel J Resolved Magnetic Resonance Spectroscopy: Simulation and In Vitro Results
AUTHOR:	Benjamin Geraghty B.Eng., (Electrical and Biomedical Engineering) University, Hamilton, Canada
SUPERVISOR:	Dr. M.D. Noseworthy
NUMBER OF PAGES:	ix, 91

To my parents...

Abstract

Localized Magnetic Resonance Spectroscopy is a non-invasive tool that offers insight into physiological status via signals arising from biological compounds. Unambiguous evaluation of said signals, however; is intrinsically limited by self interference through signal overlap. J Resolved Spectroscopy introduces an additional dimension to the measured signal which reduces overlap at the cost of increasing the scan duration. Compressed Sensing is a growing mathematical framework that asserts that under certain conditions, if a signal admits a sparse representation then it can be recovered from fewer measurements than required by classical signal theory. This framework has been successfully applied in high resolution Nuclear Magnetic Resonance experiments, justifying the investigation into its applicability in the realm of localized Magnetic Resonance Spectroscopy. The problem is addressed by optimizing the Compressed Sensing recovery on model spectra and evaluated in vitro through a parametric approach.

Acknowledgements

I wish to thank my supervisor professor Michael D. Noseworthy, Ph.D for giving me the opportunity to work in his lab. Thank you to Norm Konyer, who always had an open door policy for bouncing ideas off. Thank you to Dr. Albert Thomas who spared many hours of his time to lend his expertise and to listen to my thoughts and questions. Thank you to Dr. Krzysztof Kazimierczuk, an NMR physicist from Poland who, depsite having nothing to gain from our correspondence, always answered my questions with precision and thoughtfulness. Thank you to Dr. Mathews Jacob for letting me share my humble thoughts on topics in your domain. I would also like to thank Dr. Tim Davidson for sharing insight on some very simple ideas without ever making me feel dim, as those moments of clarity still resonante with me. I'd also like to acknowledge everyone in the Imaging Research Center who have shared a laugh or a tear along the way.

Thank you to my parents, to whom I owe everything.

Thank you to May, for her patience.

Contents

List of Figures

Chapter 1

NMR Physics

1.1 Spin and the Spin Dipole Moment

Spin is an illusive but fundamental property of matter that must be taken into account when considering a complete description of a quantum mechanical system. To this end, it may be implied that spin somehow contributes to the combination of discrete states that said system may adopt. For the following discussion, a theoretical ensemble of isolated protons will be considered without regards to their internal composition and accepting that the quantum spin number, $I = \frac{1}{2}$. owes to their underlying nature. The difficulty with fully grasping the concept of spin can be attributed to the lack of an everyday macroscopic manifestation of a protons spin state. This is in stark contrast to the idea of mass, charge and magnetism which all originate on a microscopic scale and give rise to the familiar physical laws that govern the tendencies of macroscopic bodies. As alluded to, in the absence of cleverly devised external agencies, the manifestation of spin from individual particles cancel out on macroscopic scales, precluding the existence any familiar phenomena that may be

attributed to spin.

Pragmatically speaking, the most important property of spin is that it is closely linked with the magnetic dipole moment of a particle. It is this property that allowed the early experimentalists to quantify the quantum nature of protons. Just as the rotational angular momentum vector of a rigid body about a symmetry axis has an associated moment, the intrinsic angular momentum vector, or spin, of a proton has an associated magnetic dipole moment, μ. The relationship is given by:

$$\overrightarrow{\mu} = \gamma\hbar\overrightarrow{I} \tag{1.1}$$

Where $\hbar I$ is the angular momentum vector and γ is a nuclei specific constant in units sT^{-1}, denoted the gyromagnetic ratio. Dimensional analysis suggests that if the magnetic dipole moment is expressed in units ofJT^{-1} and the reduced Planck constant in units of $J \cdot s$, then the gyromagnetic constant serves to equalize the equation. This leaves spin, I, as a dimensionless quantity that alone specifies the direction of the magnet moment vector in space. Note that the nuclear specific gyromagnetic ratio can be either positive or negative, in each case determining the parallel or antiparallel alignment of the angular momentum vector with respect to the magnetic moment vector.

Classical field theory states that the potential energy of a magnet dipole in the presence of a static external field is dependent on the magnitude of the interacting fields as well as their relative orientations. This phenomenon can be expressed as dot product operation:

$$E = -\overrightarrow{\mu} \cdot \overrightarrow{B} \tag{1.2}$$

where $\overrightarrow{B}$ is the applied magnetic field. For a static uniform magnetic field applied along the +z direction in a right handed Cartesian space, the field is denoted $\overrightarrow{B} = B_0\hat{z}$. The equivalent potential energy due to the interaction of a magnetic dipole with such a field is given by:

$$E = -\overrightarrow{\mu} \cdot \overrightarrow{B} = -\mu_z B_0 \tag{1.3}$$

Substitution of equation 1.1 into 1.3 leads to the result:

$$E = -\mu_z B_0 = -\gamma\hbar I_z B_0 \tag{1.4}$$

The quantity I_z represents the z-directed or longitudinal component of the intrinsic angular momentum of a particle. Analogous to the discrete vibrational modes of a plucked violin string, quantum mechanics describes the quantization of spin into discrete states. The allowed values of I_z are determined by the magnetic quantum number mI which takes integer values $mI = [I, I-1, \ ... \ , -I]$. In the case of protons with $I = \frac{1}{2}$, the allowed values of mI are $m = [-\frac{1}{2}, \frac{1}{2}]$. The consequence for isolated spin-$\frac{1}{2}$ particles in the presence of an applied magnetic field is that through a phenomenon known as the Zeeman effect, two distinct energy levels of the magnetic dipole moment are observed, namely:

$$E_{+\frac{1}{2}} = -\gamma\hbar(m_{+\frac{1}{2}})B_0 = -\frac{\gamma\hbar B_0}{2}$$

and

$$E_{-\frac{1}{2}} = -\gamma\hbar(m_{-\frac{1}{2}})B_0 = \frac{\gamma\hbar B_0}{2}$$

The energy separation of these two discrete states, ΔE, is directly proportional to

the magnitude of the applied field and can be calculated as:

$$\Delta E = E_{-\frac{1}{2}} - E_{+\frac{1}{2}} = (\frac{\gamma\hbar B_0}{2}) - (-\frac{\gamma\hbar B_0}{2}) = \gamma\hbar B_0 \tag{1.5}$$

This energy separation, which is linearly dependant on the external field strength, determines the energy required for a transition between the two states. The two states are commonly referred to as being either up/down or parallel/antiparallel, a convention that references the relationship of the orientation of the dipole moment to the external field. In the absence of magnetic fields beyond the static external field, the two allowed transitions between the higher and lower energy states are equiprobable and the net magnetization of the sample is cancelled out.

1.2 Spin Populations

In the liquid state, compounds undergo vigorous rotational and translational motion by virtue of the thermal energy present in the sample. The consequence for an isolated spin system is that it is subject to not only the static external field, but also time varying magnetic fields arising from the electrons of compounds comprising the microscopic environment. The two contributions of the time varying fields are due to the orbital motion of electrons which in turn generate magnetic fields as well as the magnetic moments of said electrons. Through a process known as spin-lattice relaxation, the dipole moment of a single spin is influenced by the essentially random perturbations of fields arising from its environment, causing it to stray from its initial orientation.

The effect of the random perturbations is that the spin system arrives at a state

of thermal equilibrium with its surroundings. The equilibrium energy levels are populated according to a Boltzmann distribution:

$$\frac{N_+}{N_-} = e^{\frac{-\Delta E}{kT}} \tag{1.6}$$

The magnitude of ΔE for a given spin system is bounded by the external magnetic field strength, B_0. For realizable field strengths at ambient temperatures, the energy separation ΔE is much smaller than the quantity kT, where k is the Boltzmann constant and T is the temperature (in Kelvins). The function e^x can be written as a Taylor expansion of the form:

$$e^x = 1 + \frac{x}{1!} + \frac{x^2}{2!} + \frac{x^3}{3!} + ... \quad -\infty < x < \infty \tag{1.7}$$

Since $\Delta E << kT$, the population distribution can be well approximated to the first order, resulting in the form:

$$\frac{N_+}{N_-} = 1 - \frac{\Delta E}{kT} \tag{1.8}$$

The utility of this form follows from interpretation: the population differences between the allowed states are very small and the lower energy state slightly favoured at thermal equilibrium. Since the dipole moments of an ensemble of spins distributed among the two states largely cancel out, the consequence is that the net magnetization along the direction of the external magnetic field is very weak.

1.3 The Classical Description of Precession

Above all, scientists and engineers are pragmatists, suggesting that if various treatments provide the same predictions, then the simplest description will suffice. This is analogous to favouring Newtonian physics over Einsteins Relativity to describe projectile motion on Earth. Although Relativity may provide a more complete theoretical framework, the tradeoffs of theoretical complexity against insight must be evaluated. This is certainly the case when attempting to describe the phenomena associated with magnetic resonance within the scope of this thesis and for this reason; the simplified classical picture is favoured over the obtuse, but complete quantum mechanical description. Just as Newtons theories begin to break down at interstellar distances, requiring Relativity to make accurate predictions, the classical description of spin systems also has its limitations; however, the cases in which the classical description can no longer produce accurate predictions are well outside the scope of this thesis.

Thus far, only the z-directed component of the magnetic dipole moment has been considered with no treatment of the transverse or x and y components. Although all components of the dipole moment are subject to the z-directed external field as well as the time-varying fields arising from the microscopic environment, the establishment of thermal equilibrium encompasses only energy considerations with respect to the potential energy arising from the external field. By virtue of the dot product, the potential energy of the dipole moment in the external field is independent of the transverse components; therefore, for a single spin, there is no favourable transverse orientation. Combining the polarization effect along the external field with this conclusion, an ensemble of spins will on average trace out a cone oriented along and

against the field. In order to arrive at a phenomenological description of the dynamic interaction between spins and magnetic fields, a single spin on this cone is first considered. A single spin polarized along the external field with an arbitrary transverse

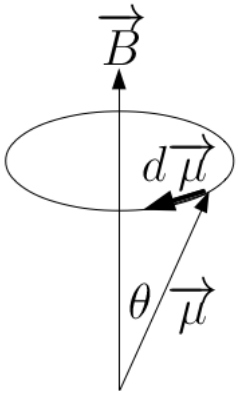

Figure 1.1: Spin Precession

component will adopt an angle between the magnetic moment and the external field denoted as θ. A classical treatment suggests that the dipole will experience a torque, τ, due to the system, which can be expressed conventionally as a cross product, citing the equivalence between the torque and the rate of change of angular momentum:

$$\overrightarrow{\tau} = \frac{d\overrightarrow{L}}{dt} = \hbar\frac{d\overrightarrow{I}}{dt} = \overrightarrow{\mu} \times \overrightarrow{B}_0 \tag{1.9}$$

Substitution of equation 1.1 into 1.9 gives:

$$\frac{d\overrightarrow{\mu}}{dt} = \gamma\overrightarrow{\mu} \times \overrightarrow{B}_0 \tag{1.10}$$

Equation 1.10 describes a first-order differential vector equation that has the component wise solution:

$$\overrightarrow{\mu}(t) = \mu_x(t)\hat{x} + \mu_y(t)\hat{y} + \mu_z(t)\hat{z} \tag{1.11}$$

where

$$\mu_x(t) = \mu_x(0)cos(\gamma B_0 t) + \mu_y(0)sin(\gamma B_0 t)$$

$$\mu_y(t) = \mu_y(0)cos(\gamma B_0 t) - \mu_x(0)sin(\gamma B_0 t)$$

$$\mu_z(t) = \mu_z(0) \tag{1.12}$$

Examination of the cross product suggests that since the external field is aligned along the z-direction, only the components of the dipole moment in the transverse plane will experience a torque. This is the cause of the clockwise precession about the field.

Equations 1.12 describes the precessional motion of a spin within an external magnetic field at angular frequency γB_0. Comparing with equation 1.5 which describes the energy separation between the two longitudinal spin states and citing the Planck equation, $E = \hbar\omega$:

$$\Delta E = \gamma \hbar B_0 \qquad \rightarrow \qquad \omega_0 = \gamma B_0 \tag{1.13}$$

Combining constants in this way arrives at the Larmor equation which connects the species specific gyromagnetic ratio, γ, and the magnitude of the applied field, B_0, to the precessional frequency and is a fundamental identity used throughout magnetic resonance theory. The important note to be made about the behaviour of spins in a magnetic field is that they precess around the field at the Larmor frequency, a frequency which scales linearly with the magnitude of the applied field.

1.4 Magnetization and the Bloch Equations

Rather than considering the magnetic moments due to a set of spins explicitly, it is convenient to describe the magnetization vector, $\overrightarrow{M}$, arising from the summation of all the dipole moments of the spins contained in a sample volume, V. This simplification assumes a volume small enough such that small, spatially varying deviations in the static field can be ignored, but large enough that it contains an appreciable number of spins. Such a volume of spins is termed an isochromat, aptly named given the assumption that all the spins contained precess at the same frequency. The magnetization vector stated formally is then:

$$\overrightarrow{M} = \frac{1}{V} = \sum_{Spins} \overrightarrow{\mu}_i \tag{1.14}$$

The equation of motion describing the response of the sample magnetization can be obtained by substituting the equation above into equation 1.10 and observing the linearity of the differential operator and cross product, giving:

$$\frac{d\overrightarrow{M}}{dt} = \gamma \overrightarrow{M} \times \overrightarrow{B} \tag{1.15}$$

Equation 1.15 provides the grounds for deriving a single differential vector equation that simultaneously describes the response of the magnetization to external magnetic fields as well as the interactions between the spins through relaxation processes. The following phenomenological equation is known as the Bloch equation:

$$\frac{d\overrightarrow{M}}{dt} = \gamma \overrightarrow{M} \times \overrightarrow{B} + \frac{1}{T_1}(M_0 - M_z)\hat{z} - \frac{1}{T_2}\overrightarrow{M}_{\perp} \tag{1.16}$$

By inspection, as the constants T_1 and T_2 approach infinity, the spins are said to be non-interacting and the Bloch equation simplifies to the form derived in equation 1.15. The second term of the Bloch equation accounts for the process alluded to in the prior discussion about the establishment of thermal equilibrium and is known as spin-lattice, or longitudinal relaxation. The solution of the Bloch equation in the presence of a static external magnetic field along the z direction is given by:

$$M_z(t) = M_z(0)e^{-\frac{t}{T_1}} + M_0(1 - e^{-\frac{t}{T_1}}) \tag{1.17}$$

where M_0 is the equilibrium value of the z-directed or longitudinal magnetization. The experimentally measured time constant, T_1, plays a key role in determining how long to wait between successive NMR experiments and is generally on the order of 100-1000 msec for compounds of interest in the human body.

1.5 Transverse Magnetization

Before examining the role of the role of T_2 in the third term of the Bloch equation, it is instructive to set the stage for when this decay process will occur. Consider an applied field that is the superposition of the static, z-directed field B_0 as well as a circularly polarized field $\overrightarrow{B}_1(t) = B_1[cos(\omega t)\hat{x} + sin(\omega t)\hat{y}]$. In a frame of reference rotating at the same precessional frequency as the spinning dipoles, the field can be written as $\overrightarrow{B}_1(t) = B_1\hat{x}'$ where the primed unit vector denotes the spinning frame. Notice that in this frame, the ωt dependence of the polarized field disappears.

This convention is analogous to solving a simple projectile motion problem observed from the surface of the earth rather than observed from orbit. From orbit,

the parabolic trajectory is complicated by the rotation of the Earth and is simplified by choosing a more convenient frame of reference. The solution to equation 1.15 for $\overrightarrow{B} = B0\hat{z}' + B_1\hat{x}'$, assuming the magnetization is initially at thermal equilibrium i.e. $M_x(0) = My(0) = 0$ and $M_z(0) = M_0$, is given in the rotating frame as:

$$M_{y'}(t) = M_0 sin\phi(t), \qquad M_{z'}(t) = M_0 cos\phi(t) \tag{1.18}$$

where the phase argument is:

$$\phi(t) = \int_{t_0}^{t_f} \gamma B_1(\tau) d\tau \tag{1.19}$$

Equation 1.19 is the flip equation which is comprised of two main elements, time duration and magnitude. The circularly polarized field, B_1, is termed an RF pulse, due to its radio-frequency oscillation and the fact that it is of finite duration. The generation and reception of RF fields is not covered in this text, but the salient point is that it is the combination of a static external field with an RF pulse that allows the magnetization vector to be rotated into the transverse plane from which its presence can be measured as an induced voltage.

Holding the amplitude constant and modifying the pulse duration is the simplest way an RF pulse can be designed to allow an experimentalist to control the trajectory of the magnetization vector. The two most commonly used RF pulses are excitation and refocusing pulses. An excitation pulse is designed such that the angle subtended by the magnetization is 90° , which from equation 1.18, fully rotates the equilibrium magnetization, M_0, into the transverse plane.

After the excitation pulse has been turned off, the magnetization vector freely

precesses in the transverse plane about the external field. Reintroducing the Bloch equation, if the time constant $T_2 \rightarrow \infty$, then the transverse magnetization in a stationary frame is described for all $t > 0$ as:

$$\overrightarrow{M}_{\perp}(t) = \overrightarrow{M}_{\perp}(0)e^{-i\omega_0 t} \tag{1.20}$$

where complex notation is used to describe the transverse magnetization:

$$\overrightarrow{M}_{\perp}(t) = \overrightarrow{M}_{x}(t) + i\overrightarrow{M}_{y}(t) \tag{1.21}$$

1.6 Transverse Relaxation

In reality, the individual dipole moments that comprise the magnetization vector are all subject to slightly different macroscopic environments that necessarily contain different field perturbations. Since these perturbations serve to add and subtract from the external field applied to the spins, each dipole precesses at minutely different frequencies. Over the course of time, deviations in precessional frequencies degrade what is said to be the phase coherence of the magnetization, which reduces the magnitude of the transverse magnetization vector. This process can be modelled in its simplest form as a mono-exponential decay, evidenced by the solution of the Bloch equation in the rotating frame:

$$\overrightarrow{M}_{\perp}(t) = \overrightarrow{M}_{\perp}(0)e^{-\frac{t}{T_2}} \tag{1.22}$$

The process of transverse relaxation reflects signal losses due to the tendency of the system towards entropy, and is irrecoverable. T_2 losses place fundamental limitations on the duration of a single experiment and limit how many pulses of finite duration may be played before the phase coherence of the magnetization is lost entirely. For pure water the time constant T_2 is approximately of the same order as T_1. In biological systems this near equality diminishes leaving T_1 being on the order of 10 to 100 times that of T_2.

In practice, the true signal decay rate is even faster than suggested by the experimental relaxation rate, T_2, due to the unavoidable but minimizable inhomogeneity of the external magnetic field. Small variations in the field lead to a separate decay rate, T_2^* , when combined in the following way:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{1.23}$$

provides the true decay rate, T_2^*. The $(T_2')^{-1}$ term is often replaced by $\gamma \Delta B_0$, reflecting the field inhomogeneity term. In contrast to the intrinsic process of transverse relaxation, losses due to field inhomogeneity can be recovered through the use of refocusing pulses. As individual dipole moments lose coherence due to field variations, which can be viewed as fanning out in the rotating frame. The spins experiencing slightly higher fields lead the reference frame while the spins experiencing slightly lower fields lag behind. By rotating this fan of spins within the transverse plane by 180° through the use of a refocusing RF pulse, losses due to field variations can be recovered.

A common example used to understand this phenomenon is the runners on a track analogy. Consider three runners, each one faster than the next, who are competing

in a race. In a traditional foot race, the fastest runners would always cover the predetermined distance in the shortest time. If instead the runners are asked to run an arbitrary amount of time, and then at a single instant are instructed to turn around and run back to the start, they would all return to the starting line at exactly the same finishing time, a duration twice as long as the time elapsed before the instruction to return to the start was given. This occurs because the fastest runner must travel the largest distance back and the slowest runner the shortest. In the language of magnetic resonance spectroscopy, this phenomenon is known as a spin-echo and will be treated in more detail in section 2.3.

1.7 Chemical Shift and J Coupling

The true utility of magnetic resonance has been exploited through the past several decades by analytical chemists who have made great use of magnetic resonance experiments in order to elucidate the chemical structure of complex compounds. It is difficult to imagine that the application of magnetic resonance theory would have matured to its present state if the interactions between spins and magnetic fields were limited to the basic concepts discussed thus far in fact, it is the deviations from the theory presented above that have motivated the majority of the development of magnetic resonance.

Consider the case discussed above of a sample of isolated spin-$\frac{1}{2}$ particles in the presence of an inhomogeneous external field. It was argued that each spin isochromat experienced a slightly different magnetic field and that through the Larmor equation, it is surmised that each set of spins would precess at a slightly different frequency resulting in a spread of frequencies across an entire sample. For the following discussion,

the scenario is modified such that the external field again is considered completely uniform and the concept of isolated spins is abandoned in favour of the more realistic situation in which the spin-$\frac{1}{2}$ nuclei are chemically bound into molecules, collections of atoms that are electrochemically linked.

As before, the field is defined as being completely homogenous on a macroscopic scale; however, on a microscopic scale, this assumption breaks down. As an illustration of this phenomenon, consider the simple and abundant compound: water. Water, as evidenced by its molecular formula H_2O, contains two hydrogen atoms covalently bound to an oxygen atom, contributing to the classification of its molecular shape as bent. A generalized interpretation of such a configuration is that the atoms com-

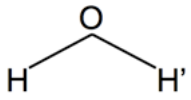

Figure 1.2: Water

bine by sharing their electrons, producing an electron cloud of a well defined average spatial distribution. As discussed previously, the orbital motion of charged electrons generate their own local, microscopic magnetic fields. The result for magnetic resonance is that the spin-$\frac{1}{2}$ hydrogen nuclei experience the superposition of the external field and the electronically induced local magnetic field. This phenomenon is termed the chemical shift, and can be expressed as a deviation from the isolated Larmor frequency through a shielding parameter, σ:

$$\omega_0 = \gamma(1 - \sigma)B_0 \tag{1.24}$$

To good approximation, the shielding parameter can be regarded as a constant, independent of the applied magnetic field. This implies that just as the Larmor frequency for an isolated spin scales linearly with frequency, so does chemical shift. Not only does increasing the magnitude of the applied field increase the degree of polarization that leads to increased signal, chemical shift dispersion is also increased, leading to increased spectral resolution.

Due to the symmetry of the water molecule, both hydrogen nuclei are subject to the same shielding effect generated by the electron cloud and are said to be chemically equivalent. Chemically equivalent nuclei all share an identical chemical shift and will possess the same resonant frequency for a given field strength. Since both the chemical shift and Larmor frequency scale with the applied field, it is convenient to quote the chemical shift, δ, with respect to a reference compound at the same field strength as the experiment. This dimensionless metric is field-independent and given as:

$$\delta = \frac{\omega^0 - \omega^0_{ref}}{\omega^0_{ref}} \tag{1.25}$$

Chemical shift is one of the fundamental phenomena of magnetic resonance that make it a useful tool for identifying the chemical structure of an unknown substance. In this way, spins can act as if they were radio frequency transmitters, sending out a signal with that is encoded with their unique chemical fingerprint.

The concept of chemical equivalence is in contrast to the concept of magnetic equivalence. Due to the symmetry of the water molecule, both hydrogen nuclei are magnetically and chemically equivalent. In order to illustrate the concept of magnetic equivalence, consider another simple compound, 1-1-difluroethene, CH_2F_2 (figure 1.3). The hydrogen nuclei in 1-1-difluroethene are chemically equivalent due to

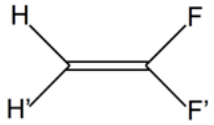

Figure 1.3: 1-1-Difluroethene

symmetry considerations, this is to say if H' were exchanged with H, the chemical nature of the compound is unchanged and by this line of reasoning, $\delta_{H'} = \delta_H$. Although both hydrogen nuclei are chemically equivalent, due to the structure of 1-1-difluroethene, they are not magnetically equivalent. The magnetic coupling between H and F is not equivalent to that between H' and F, similarly for H/H' to F'. This results in a non-zero coupling interaction known as scalar or J coupling.

J coupling is a physical phenomenon that acts through the electrons in chemical bonds and contributes to the possible states which the spins comprising a system may adopt. Just as the magnetic dipole arising from the hydrogen nuclei will adopt either an orientation parallel or antiparallel to the applied field, so too will the moments arising from the bonding electrons. In isolation, the electron associated with the hydrogen nuclei will preferentially adopt an antiparallel orientation with respect to its hydrogen nuclei as the total energy is minimized. In the case of electrons shared in a chemical bond, Paulis exclusion principle states that two electrons (fermions) may not have identical quantum states, necessitating the antiparallel (up/down) orientation between the bonding electrons. This means that for one of the nuclei who's atom is sharing a bonding electron, it is forced into a higher energy up/up state with respect to the electron, increasing the total energy of the spin system.

The situation is summarized in table 1.1 for the simple case of spin I and S, with Larmor frequencies in Hz, υ_I , υ_s, where α and β denote the spin state being parallel

and antiparallel respectively. Energy is expressed in frequency units by dropping Planks constant.

I S	Energy (uncoupled)	Energy (coupled)
$\beta\beta$	$E_4 = \frac{1}{2}v_I + \frac{1}{2}v_S$	$E_4 = \frac{1}{2}v_I + \frac{1}{2}v_S + \frac{1}{4}J_{IS}$
$\beta\alpha$	$E_3 = \frac{1}{2}v_I - \frac{1}{2}v_S$	$E_3 = \frac{1}{2}v_I + \frac{1}{2}v_S - \frac{1}{4}J_{IS}$
$\alpha\beta$	$E_2 = -\frac{1}{2}v_I + \frac{1}{2}v_S$	$E_2 = -\frac{1}{2}v_I + \frac{1}{2}v_S - \frac{1}{4}J_{IS}$
$\alpha\alpha$	$E_1 = -\frac{1}{2}v_I - \frac{1}{2}v_S$	$E_1 = -\frac{1}{2}v_I - \frac{1}{2}v_S + \frac{1}{4}J_{IS}$

Table 1.1: Spin-State Energy Levels

Note how the sign of J_{IS} depends on the spin states of both nuclei. The following tables summarize the allowed ($\Delta m = \pm$) transitions associated with the four possible spin states of the IS system:

I Transitions	ΔE	Frequency
$\beta\alpha \rightarrow \alpha\alpha$	$E_3 - E_1$	$v_I - \frac{1}{2}J_{IS}$
$\beta\beta \rightarrow \alpha\beta$	$E_4 - E_2$	$v_I + \frac{1}{2}J_{IS}$

Table 1.2: I Spin Transitions

S Transitions	ΔE	Frequency
$\alpha\beta \rightarrow \alpha\alpha$	$E_2 - E_1$	$v_S - \frac{1}{2}J_{IS}$
$\beta\beta \rightarrow \beta\alpha$	$E_4 - E_3$	$v_S + \frac{1}{2}J_{IS}$

Table 1.3: S Spin Transitions

The phenomenon of J coupling is a reciprocal, internal mechanism by which neighbouring spins can sense each others presence through the electrons of a chemical bond. Note that in the absence of J coupling, a situation that arises when the spins are said to be magnetically equivalent, the two transitions for each spin have equal energy and they are said to be degenerate. Breaking of the degeneracy results in frequency splitting and can be exploited to elucidate the connectivity of atoms within a compound

of interest. It is important to note that unlike chemical shift, J coupling is field independent since it arises through a mechanism internal to the compound rather than through the degree of magnetic shielding of an external field.

Chapter 2

Localized J Resolved Spectroscopy

Magnetic Resonance Spectroscopy (MRS) permits non-invasive observation of NMR signals arising from biological systems. Although the focus of research and clinical endeavours involving MRS have focused on proton spectroscopy of the brain, the technique can be generally utilized with other MR visible species such as 31-P, 13-C and 23-Na arising from all areas of the body. The following sections describe the general application of localized proton spectroscopy and the extension to 2D J Resolved spectroscopy.

2.1 Localized Spectroscopy and PRESS

The term localized spectroscopy implies the use of an MR scanner to acquire line spectra from a selective region. The most primitive technique for localization is termed "pulse acquire". Pulse acquire methods involve the application of a 90° RF pulse to excite the spins directly followed by signal reception, generally performed using a surface coil. Since the receptive field of a surface coil is limited by its size, the

signal is acquired from a region proportional to the radius of the coil. Localization is a function of the coil geometry such that signals measured in pulse acquire experiments are an aggregate of all signals within the sensitive region of the coil.

A more sophisticated localization approach is to discriminate signals from a Region-Of-Interest (ROI) by the use of Single Voxel (SV) methods. This thesis is concerned only with the SV technique known as Point REsolved SpectroScopy (PRESS), but there exist many implementations for confining the final signal within a physical ROI. For a comprehensive review of localization methods that have been applied for the acquisition of MRS data, see [1].

The PRESS methodology involves the application of tailored RF pulses in concert with linear magnetic field gradients. The Larmor equation (1.13) relates the precessional frequency to the magnitude of the applied field. Ignoring variations due to the chemical shift, the introduction of a linear magnetic field gradient induces a spatial dependance upon the total field such that the position dependant Larmor frequency is given by:

$$f(x) = \frac{\gamma}{2\pi}(B_0 + G_x \cdot x) \tag{2.1}$$

where G_x is the linear gradient in units $T \cdot m^{-}1$. The following figures illustrate this phenomenon graphically:

As the frequency depends linearly on the applied field, and the field depends linearly on position, then the Larmor frequency depends linearly on position. Section 1.? discussed the resonance condition necessary for flipping spins into the transverse plane. Since the Larmor frequencies of spins subject to a spatially varying field are also spatially varying, the frequency response of the RF pulse must correspond with a square-shaped bandwidth. Consider the following figure depicting a well known

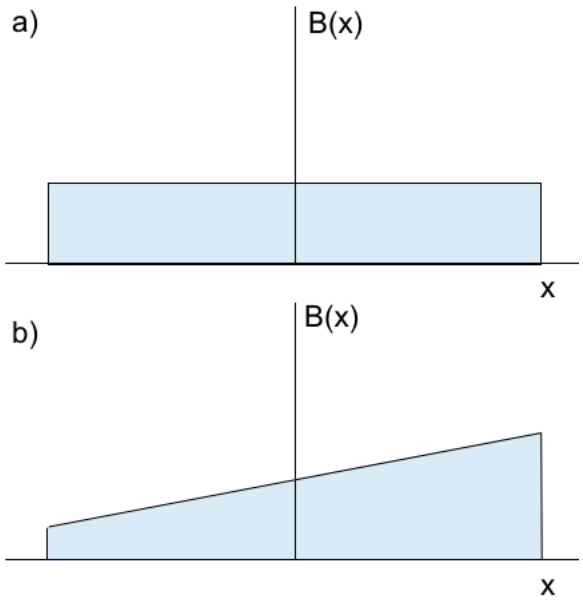

Figure 2.1: Magnetic Field vs Position a) Static Field b) Addition of Linear Gradient

Fourier Transform pair:

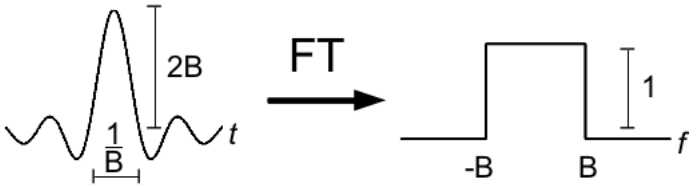

Figure 2.2: Sinc Filter

The frequency response of the ideal sinc waveform is a rectangular function with bandwidth B. Let $s(t)$ denote the sinc filter,

$$s(t) = 2Bsinc(2Bt) = 2B\frac{sin(2\pi Bt)}{2\pi Bt} \tag{2.2}$$

then $S(f)$ is the Fourier Transform of $s(t)$, given in conventional frequency as:

$$S(f) = rect(\frac{f}{2B}) = \begin{cases} 1, & \mid f \mid \leq B \\ 0, & \mid f \mid > B \end{cases} \tag{2.3}$$

The application of a sinc shaped RF wave in concert with a linear gradient allows the excitation of spins confined to a slab of thickness determined by the slope of the gradient. This is illustrated in figure 2.3 where through the linear gradient, frequency is related to position. The slice thickness is flexibly controlled by manipulating the magnitude of the gradient rather than by changing the bandwidth of the sinc wave.

In real-world applications, the RF waveform must have a finite duration. Truncation of the waveform leads to Gibb's phenomenon where ripples are seen in the pass band. Truncation artifacts can be reduced by multiplying the RF pulse by a window function. The use of numerical methods for the design of RF pulses produces more optimal pass band characteristics and these types of RF pulses are commonly employed [2]. These designs are subject to both hardware and biological constraints.

PRESS employs the application of Slice Selective (SS) excitation and refocusing pulses in three mutually orthogonal directions in order to localize a three dimensional ROI. In this sense, PRESS is considered a single-shot technique such that the final echo arising from the ROI is generated from a single pass of RF pulses and gradients.The first step is in the sequence is the application of a SS excitation pulse along a single direction. This has the result of producing transverse magnetization within a restricted slab (figure 2.4a). Subsequently, a large magnitude gradient is applied which serves to destroy the phase coherence in the transverse plane. The gradient is followed by a SS refocusing pulse which inverts the transverse orientation of excited

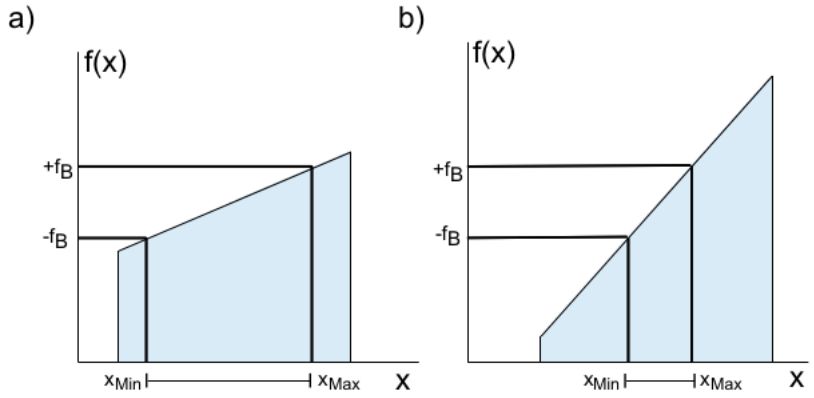

Figure 2.3: Frequency vs Position a) Weak Gradient b) Strong Gradient

spins along the next orthogonal direction. By playing out a linear gradient which is a replicate of the one used to destroy the transverse magnetization, the effects are cancelled out, leaving a column of spins along the intersection of the two orthogonal slabs (figure 2.4b). Finally, the same sequence of spoiler gradient, slice selective refocusing and spoiler gradient are played in the remaining orthogonal direction, restricting the excited signal to the intersection of the three orthogonal slabs (figure 2.4c).

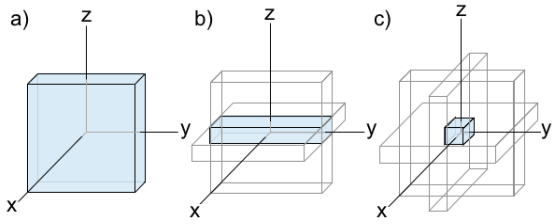

Figure 2.4: Visualization of PRESS Excitation a) SS excitation along x, b) SS refocussing along z and c) SS refocussing along y

To summarize the PRESS methodology, the final echo is confined to the defined ROI such that spins outside the intersection of the three planes are either not excited or not refocussed. Without the action of the spoiler gradients, there would exist residual transverse magnetization arising from outside of the ROI (the entire slab

of figure 2.4a). Although it would be subject to greater signal losses through T_2^* decay, it would still significantly contribute to the final signal. By destroying the transverse phase coherence and selectively recovering it within an orthogonal plane, the signals arising from outside the ROI are minimized thereby achieving a sharp excitation profile. Theoretically speaking, the ordering of the gradient directions plays no bearing on the quality of volume localization, although specific vendors may recommend a particular ordering based on hardware specifications.

The final major ingredient of any PRESS pulse sequence is the application of global water suppression. Since the concentration of water or in tissue is greater than concentration of metabolites by four orders of magnitude [1] it is necessary to attempt to reduce the solvent signal as much as possible before acquisition. Although it is possible to reduce the appearance of water in the spectrum through post-processing, the dynamic range of the ADC is finite such that an unsuppressed water peak will saturate the digital signal, leaving only a fraction of the bits to account for the metabolite signals. There exist several implementations of water suppression whose treatment will not be considered in this thesis. For a comprehensive overview of various techniques refer to [3] and the references within. Regardless of the specific water suppression module, the goals are always the same: eliminate as much as the solvent signal as possible without distorting the spectral signals of interest. The combination of water suppression and volume localization can be summarized in a schematic of the Pulse Sequence Design (PSD) for PRESS (figure 2.5).

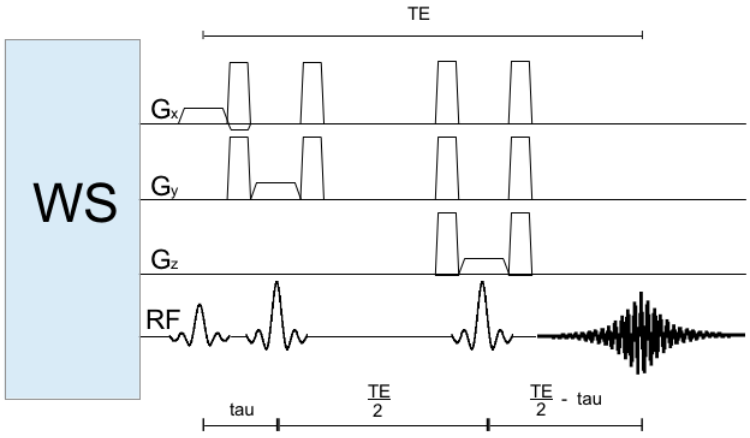

Figure 2.5: PRESS Pulse Sequence

2.2 MRS Acquisition Parameters

The operator is free to define various parameters for the an MRS experiment, all of which are associated with various tradeoffs. The Spectral Width (SW) defines twice the bandwidth of the signal and is conventionally fixed at an appropriate value for proton spectroscopy. For a given system, the ADC will acquire data at a constant rate on the order of ±250Khz. Since the signal is heavily oversampled, the final output signal is subject to electronic filters which limit the bandwidth of the signal. The resulting SW can be selected from various predefined values by the operator. The total acquisition time for a single echo is a function of the sampling or dwell time (SW^{-1}) as well as the total number of measured points, another parameter under the operators control. For a fixed SW, this allows the operator to make a choice about the tradeoff of Signal-to-Noise-Ratio (SNR) for resolution. The reason why this is a tradeoff is because as discussed in chapter 1, the NMR signal can be modelled as a sinusoidally modulated decaying exponential. Assuming white gaussian noise of a given variance, the SNR per measured point decreases over the course of a single

acquisition to the point at which the signal has completely decayed away and only noise is measured. For a fixed number of points, increasing the SW will result in a shorter data acquisition window but at the cost of resolution.

The total duration of a single acquisition including all elements of suppression, localization and acquisition is termed the repetition time (TR). The minimum TR is bounded by the total time it takes for all of the waveforms to play out as well as acquiring the signal. In practical experiments, the TR is at least 2-3 times longer than the time that the actual RF and gradient waveforms are active. The total TR takes into account the time required for allowing the reestablishment of thermal equilibrium and such is on the order of 4-5 times the T_1 of the species under investigation. For MRS in tissue, the TR is generally set to 1-5 seconds.

The number of signal averages can also be controlled by the operator in order to improve the observed SNR. The Number of EXcitations (NEX) along with the TR make the largest contributions to the total scan time. Signal averaging is performed on the assumption that the major sources of noise sum to follow a zero mean gaussian distribution. Under this assumption, averaging many signal acquisitions should reduce the standard deviation of the measured noise by the square root of the number of averages. The justification for measured noise being assumed to obey a normal distribution is that the major sources are thermal noise from the receiver coil, thermal noise from the preamplifier and thermal noise from the sample [4]. This also explains the observed gains in sensitivity when using cryogenically cooled RF coils. The total experiment time is proportional to the TR times the NEX:

$$T_{total} \propto TR \cdot NEX \tag{2.4}$$

The size and position of the ROI is also under operator control. The position depends on what the purpose of the experiment is whereas the size is subject to several constraints. In order for optimal spatial selectivity, it would be most beneficial to make the ROI exactly the size of the area under investigation. Since the signal acquired is the sum of all of the spins within the ROI, it is often assumed that the region is homogenous. In that regard, even if it weren't homogenous, there would be no way to discern from a single line spectrum what fractions of the spectrum arose from which areas. On the other hand, as the size of the voxel becomes smaller and smaller, the total number of spins contributing to the measured signal also decreases. Given that NMR is a notoriously insensitive technique, this places limitations on how small the ROI may be. Although a large ROI may produce a stronger signal, the problem of sample heterogeneity comes into play in addition to decreased homogeneity of the magnetic field across the excited region. The tradeoff between selectivity and signal results in ROI sizes on the order of $(1 \times 1 \times 1)cm^3$ to $(3 \times 3 \times 3)cm^3$.

The echo-time (TE) is an operator defined parameter that defines the delay between the initial excitation and signal acquisition. Generally speaking, the selection of TE defines the regime of the experiment. For TE $<$ 40ms, the MRS experiment is said to be of the short echo-time variety. Conversely, for TE $>$ 40ms, it is termed a long echo-time experiment. The obvious benefit of choosing a short TE is that the transverse signal decay will have had the least time to degrade the signal, leading to higher SNR spectra. Depending on what the experimenter is looking for, by choosing a long TE, certain MR species will have already largely dephased due to transverse relaxation before signal acquisition, leading to a clearer spectrum[5]. The lower bound on TE is determined, as with TR, by the minimum time required to allow the various

waveforms to be played. The role of TE in an MRS experiment is more subtle than the other parameters and will be treated in greater detail in the following section.

2.3 Spin-Echoes and J Modulation

The concept of the spin-echo was introduced with a descriptive treatment in chapter 1 and was exploited in the theoretical development of PRESS (section 2.1). It is instructive to consider a geometrical picture of the generation of a so called spin-echo. The solution of the Bloch equations are evoked implicitly in the following development.

Figure 2.6 illustrates the spin-echo mechanism for the ideal case which neglects T_2 relaxation. A frame of reference rotating at the precessional frequency of the spins about the magnetic field is used to simplify the dynamics. Magnetization initially directed along the external static field (figure 2.6a) is converted into transverse magnetization directed along +y through the action of a 90° RF pulse along the positive x direction(figure 2.6b). The phase coherence of the transverse magnetization begins to degrade (figure 2.6c) due to the inhomogeneity of the external field (section 1.?). Application of a 180° refocussing RF pulse along the positive x direction flips the dephasing spins about the axis of excitation (figure 2.6d). Finally, due to the transverse inversion, the spins experiencing slightly lower fields (darker vectors) are able to catch up to the spins experiencing slightly higher fields (lighter vectors) as the spins refocus along the negative y direction (figure 2.6e). Recall the runners on a track analogy used to explain the spin-echo phenomenon (section 1.6). The idealized spin-echo is more complex in the case of coupled spins. Recall that scalar or J coupling is a phenomenon by which the energy of a single NMR transition is split into multiple levels

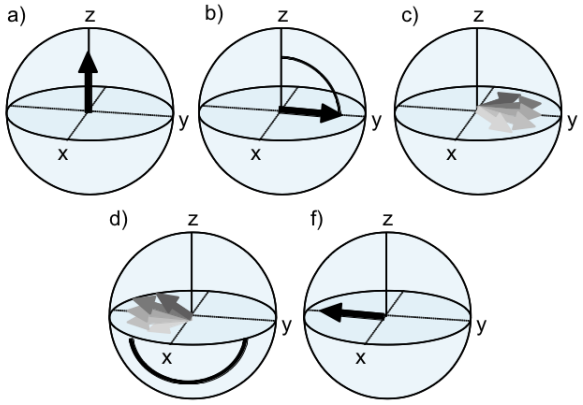

Figure 2.6: Ideal Spin-Echo a) magnetization at thermal equilibrium, b) excitation of transverse magnetization, c) phase decoherence, d) transverse inversion and e) refocussed transverse magnetization after period TE

dependant upon the orientation of nearby spins. Consider the two spin system IS. In the absence of J coupling, the two singlets resonate at their respective chemical shifts (figure 2.7a). However; if the spins are coupled, then the spectrum displays frequency splitting about the respective chemical shifts by $\pm\frac{1}{2}J_{IS}$ (figure 2.7b). The labelling of figure aids in the understanding of the nature of J coupling. Consider the I spin doublet. The I spin state for each spectral line is not specified, as each spectral line represents a transition back and forth between the higher (β) and lower (α) states. The α and β states of the S spin are reflected in the relative frequency of each I spin transition. Contributions of the S spin in the high energy β state increase the energy of one of the I spin transitions. Conversely, contributions of the S spin in the lower energy α state lower the energy of one of the I spin transitions. The same arguments apply for the appearance of the S spin transitions.Due to the interconnectedness of coupled spins, the description of the ideal spin-echo does not apply as the phenomenon known as J modulation must be accounted for. For clarity, the description of the$J$ modulated spin-echo will consider a further idealized situation

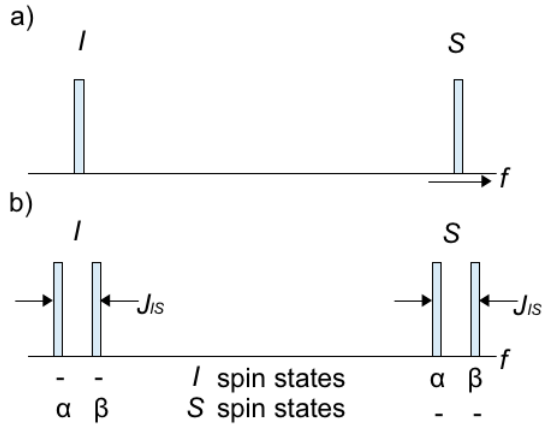

Figure 2.7: Illustration of J coupling a) uncoupled two spin system and b) coupled spin system

where both intrinsic and inhomogeneous transverse decay will be neglected. Again, the classical 90-180 spin-echo pulse sequence will be applied but this time, to a coupled two spin system. In this case, it is helpful to consider the refocussing pulse across the entire bandwidth as being a cascade of pulses, first on the I spins and then on the S spins. This treatment is equivalent in the idealized case. Assume the system

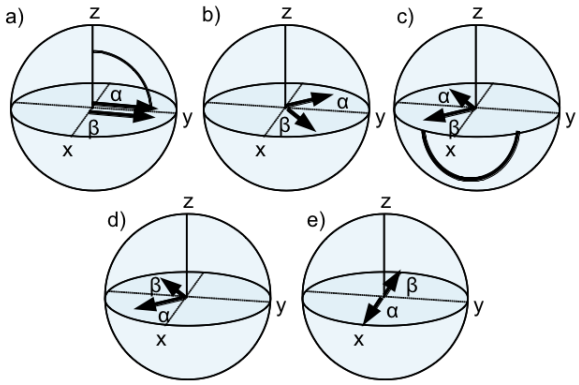

Figure 2.8: J Modulated Spin-Echo a) Excitation of α and β vectors, b) the vectors diverge due to differences in precessional frequency c) 180° pulse on I flips inverts transverse magnetization, d) 180° pulse on S swaps the α and β labels e) the α and β vectors continue precession and are not refocussed after period TE

has long established thermal equilibrium after which a 90° excitation pulse flips the z directed magnetization into the transverse plane (figure 2.8a). Consider the fate of the I spins under the influence of both the RF pulse as well as the presence of the coupled S spins. The frame of reference is rotating at the chemical shift of the I spin such that the faster β vector leads the reference frame while the slower α vector lags behind (figure 2.8b). Application of a 180° RF pulse on the I spins invert the two vectors about the $+x$ axis (figure 2.8c). Recall the physical mechanism of J coupling (section 1.?) that manifests in the relative changes in frequency at the I spin site is due to the addition of the small field generated at the S spin site. Application of a 180° RF pulse on the S spins inverts the S spin populations, consequently inverting the sense of the field at the I spin site. This has the effect of interchanging the α and β labels (figure 2.8d). The fast and slow vectors continue precessing until the end of the TE period. (figure 2.8e). To the inquisitive reader wondering how the label changing can be described in a vector model, the short answer is that it cannot. The vector description developed here requires an appeal to the spin population distribution in response to so called excitation/saturation (90°) and inversion (180°) pulses. A full quantum mechanical treatment arrives at the same result without stretching the model, albeit at the cost of complex mathematical arguments.

The scenario depicted in figure 2.8 differs fundamentally from the conventional spin-echo phenomenon due to the non-zero J coupling interaction between spins I and S. For illustration, consider the uncoupled case where the fast and slow labels of figure 2.8 could then refer to the IS singlets. Following this case, the only difference is that there would be no label swapping (figure 2.8d) and the vectors would align along the -y direction at the end of the TE period. This means that the chemical

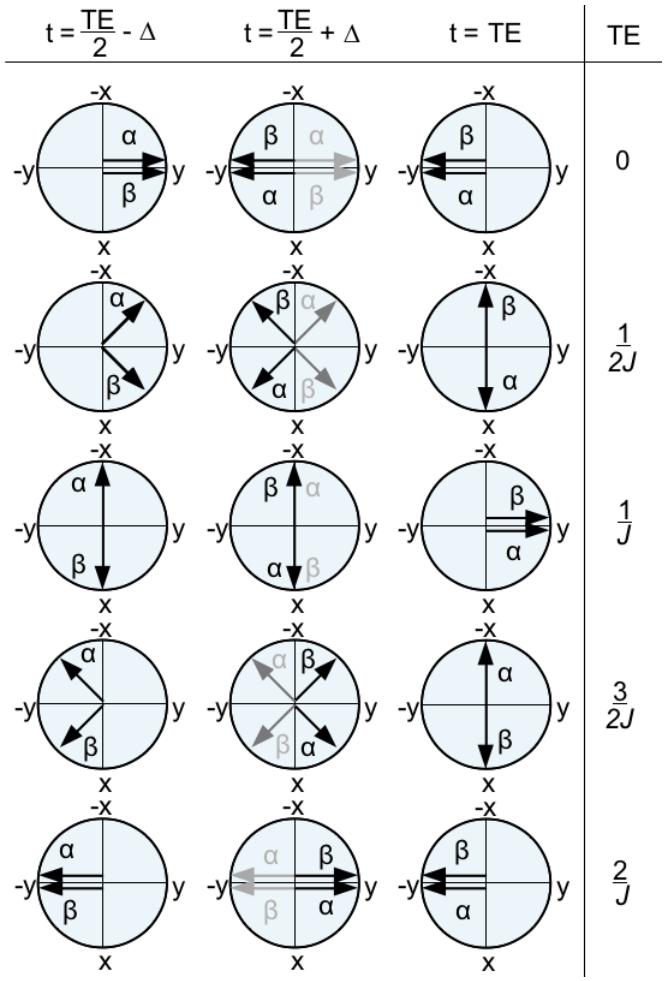

Figure 2.9: J Modulation

shift dispersion is refocussed in all spin-echo experiments, but the initial phase at the start of signal acquisition in the case of J coupled spins is dependant on the length of the TE period.

Figure 2.9 depicts schematically how the phase of the α and β vectors depend on the experimental echo-time. Again, the frame of reference is rotating at the unmodulated chemical shift of the I spin transition. The first column depicts the position of the α and β vectors after a half period of TE but directly before the application

of the 180° RF pulse. The second column depicts the position of the α and β vectors following the application of the 180° RF pulse affecting both the I and S spins simultaneously. Finally, the third column depicts the position of the α and β vectors after and additional half TE period of free precession.

2.4 2D J Resolved Spectroscopy and JPRESS

The discussion of the J modulated spin-echo revealed the periodic relationship between the phase of a spin-echo experiment and the TE through the coupling constant for a given spin system. For clarity, several simplifying assumptions were made including neglecting transverse relaxation (both intrinsic and due to field inhomogeneity) and describing only a situation involving coupling between two spins. In MRS experiments, most of the metabolites of interest are composed of dozens of coupled spins whose signals are cluttered together across a relatively narrow bandwidth. Consequently, this leads to significant signal overlap that can hinder unambiguous quantification [6] of the active metabolites. By exploiting the periodic nature of J modulation, it is possible to conceive a 2D experiment to facilitate the reduction of spectral clutter [7].

Figure 2.10d depicts the pulse sequence for the 2D J resolved experiment. The interval denoted t_1 is the so called mixing period where the spins are excited and interact. The interval denoted t_2 is the detection period at which point the signal is acquired. The time variable, t_1, corresponds to the TE of the experiment and is linearly increased over successive scans. The time variable, t_2, is simply the sampled points of the experiment hence it is commonly referred to as the direct dimension. Conversely, since for any given measured point along the t_1 dimension a new scan is

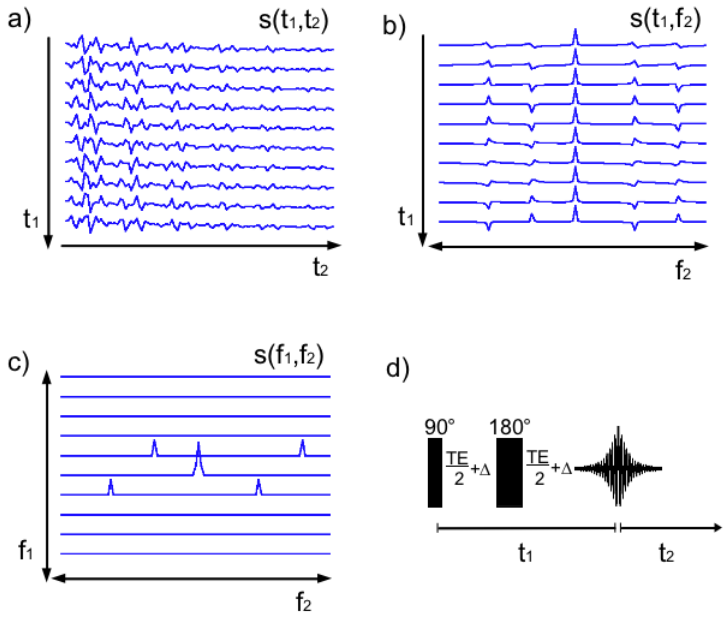

Figure 2.10: The 2D J Resolved Experiment a) An array of spin-echoes of increasing TE b) Hybrid time-frequency Spectrum c) 2D J Resolved Spectrum d) 2D J Resolved Pulse Sequence

required, it is commonly referred to as the indirect dimension. Figure 2.10a depicts an array of spin-echos, each with an incrementally larger echo time. By taking the Fourier Transform of the signal with respect to the direct dimension, an array of one dimensional spectra are revealed. Examination of the spectrum shows that there are three spins involved, two of which are coupled. The central peak is identical for all TE steps, indicating that it is not a coupled spin. On the other hand, the coupled peaks display amplitude modulation across the indirect dimension. Only the imaginary part of the spectrum has been shown to highlight the difference in phases of the amplitude modulation. By taking the Fourier Transform with respect to the indirect dimension, the 2D J resolved spectrum is revealed (figure 2.10c). The direct frequency dimension corresponds to the chemical shift of the spins while the indirect frequency dimension corresponds to frequency offsets due to J coupling. Since the

central singlet is an uncoupled spin, it lies on the $f_1 = 0$ trace of the 2D spectrum. Conversely, the coupled spins are displaced from the $f_1 = 0$ trace by $\pm\frac{1}{2}J$. The 2D J resolved experiment offers a powerful analytical tool for chemists as it allows them to decouple chemical shifts from scalar coupling, permitting the direct measurement of the magnitude and sign of coupling constants.

The spin system depicted in figure 2.10 is not complex enough to warrant a full 2D experiment, given how an experienced chemist would be able to easily deduce the chemical shifts and interactions by inspecting the 1D spectrum alone. However; often times in practical situations there exists assignment ambiguity, such as discerning from several closely spaced multiplets. In this situation, it is impossible to tell from the 1D spectrum alone which resonance line is being split due to the interaction of another line. In the realm of MRS, elucidation is not a problem as the spectral appearance of the MR visible metabolites in the body have been well established and documented[8]. The problem of MRS is not identification, rather it is quantification. Introducing a second spectral dimension can disperse problematic signal overlap and potentially improve quantification accuracy[7].

The 2D J resolved experiment can be implemented as a localized protocol by introducing an incremental delay period between the refocussing pulses of the PRESS sequence. The resulting sequence has been named J resolved PRESS and goes by the acronym JPRESS. In 2006, Shulte *et al.*. published on an improved acquisition strategy for JPRESS spectra known as maximum echo sampling[9]. The strategy was justified through mathematical and experimental evidence and is currently the conventional approach. Maximum echo sampling (figure 2.11a), in contrast with the conventional half echo sampling (figure 2.11b), offers two major advantages with a

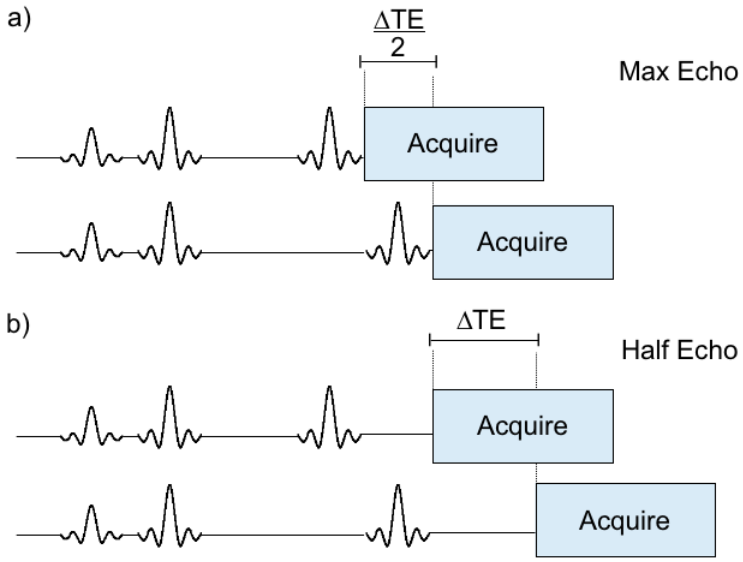

Figure 2.11: JPRESS Pulse Sequence for a) Max Echo Sampling and b) Half Echo Sampling

minor drawback of requiring minimal post processing. The first is that it improves the SNR since the signal is acquired earlier which reduces the impact of transverse relaxation. For any TE period, although the inhomogeneous broadening is refocussed, the intrinsic T_2 decay cannot be recovered. As the TE period is further incremented, the SNR of the acquisition is degraded. The second advantage of maximum echo sampling is more subtle. As alluded to, one arguable drawback is the need to apply post-processing in order to recover the conventional 2D J resolved spectrum. The reason for this is since signal detection occurs directly after the final gradient rather than after the second half TE period, then this imposes additional phase modulation upon the signal. This has the effect of modulating the signals across the indirect frequency dimension. This is corrected for by time shifting the acquired echoes, or equivalently by multiplying the hybrid time-frequency spectrum with a linear phase ramp. Refer to [9] for details. The advantageous side effect is that this tilting process

serves to deflect the solvent tail away from the metabolite spectrum. In contrast, since the residual solvent signal lies on the $f_1 = 0$Hz trace, the large sideband response in half echo sampled JPRESS can overlap with the metabolite signals when water suppression is poor. This benefit greatly outweighs the minor hinderance of applying a simple correction in post-processing. For the remainder of this thesis, the maximum echo sampling strategy will be implied in any further mention of JPRESS.

2.5 Quantification of JPRESS Spectroscopy

The primary purpose of *in vivo* MRS is to non-invasively quantify various metabolites within the body. Accurate estimations of metabolite concentrations can potentially provide clinicians and researchers with insight into physiological status [5]. The relationship between the acquired spectrum and metabolite concentrations is by no means trivial, owing to various systematic difficulties such as patient motion, non-uniform fields and inherently low SNR. Over the past several decades, numerous approaches have been developed for quantifying MRS data. For a comprehensive review comparing a wide range of approaches, refer to [10]. It has been demonstrated that in the short echo-time regime (section 2.1), quantification of many overlapping resonances is most accurately treated through the incorporation of prior knowledge as model spectra. The current gold standard in 1D MRS quantification is Linear Combination of Model Spectra (LCModel) which works in the frequency domain to estimate various parameters of a metabolite basis set in order to fit the measured data[11]. The model spectra can be either simulated through quantum mechanical numerical methods [12, 13] or measured *in vitro*, although the former is the preferred method as the latter is susceptible to instrumental errors.

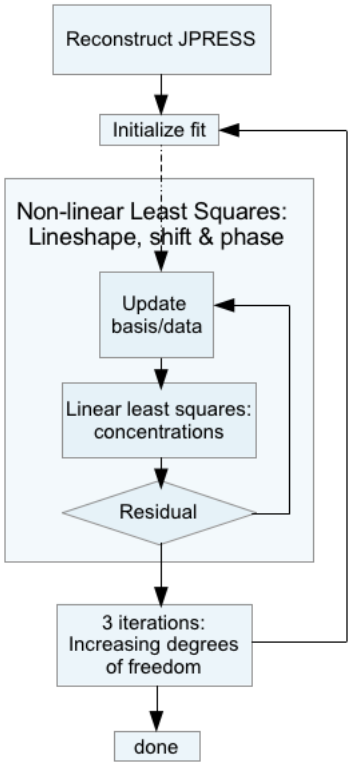

Figure 2.12: ProFit Work Flow Diagram

As discussed in section 2.2, JPRESS can be used to resolve the potentially overlapping resonances involved due to J coupling which addresses one of the technical difficulties in metabolite quantification. The least sophisticated approach to quantify J resolved spectra is to integrate the peak volumes of the magnitude spectrum. Although this approach is more or less easily implemented, it is inherently subjective and lacks accuracy in cases where peaks are not well resolved and are still overlapped. In 2006[14], Shulte *et al*. proposed to a methodology aimed to fit 2D model spectra to the measured spectrum in the frequency domain and it was dubbed Prior knowledge Fitting (ProFit). The approach is inspired by LCModel[11] for its incorporation of

prior knowledge and Variable Projection (VARPO) algorithmic scheme which combines non-linear and linear optimization [15]. Figure 2.11 depicts the work flow of the algorithm. The raw data is reconstructed into a JPRESS spectrum. The fit is initialized by phasing and frequency shifting/referencing the measured spectrum. The first pass of the algorithm only considers the dominant singlets within the spectrum in order to establish robust estimates of the lineshapes, shifts and phases. The second pass again only considers the dominant singlet resonances but incorporates a more sophisticated lineshape model to account for possible distortions in the acquisition process. The final pass considers all the remaining metabolites in the basis set. Refer to [14] for a detailed description of the various parameters considered, the use of soft constraints and the cost function used in the non-linear least squares estimation. The final concentrations are determined through linear least squares:

$$c = Re\{B^{total}\}^{\dagger} Re\{S^{r}\} \tag{2.5}$$

where S is an r dimensional vector representing the measured spectrum. The dimension r corresponds to the size of the spectral ROI used in the fit.

As mentioned in [14], the numerical burden of the non-linear fitting routine is reduced by splitting the approach into a non-linear and linear estimation scheme. The data is fit only in a narrow region of the spectrum where the metabolite signals lie, further improving numerical efficiency. The basis sets are simulated using a dedicated library [13] and are also truncated to the same region such that the Gibb's ringing in the basis set matches the measured data.

One of the most attractive features of prior knowledge methods such as LCModel and ProFit is that, under certain assumptions, the Cramer-Rao Lower Bound (CRLB)

can be computed as an objective assessment of quantification accuracy. The CRLB is a statistical lower bound for the achievable standard deviation of the estimated parameters. In 2001, Cavassila *et al.* provided a functional form of the CRLB for prior knowledge fitting of MRS data [16]. The Fisher Information Matrix is given by:

$$F = \frac{1}{\sigma^2} Re\{B^{total}\}^T Re\{B^{total}\} \tag{2.6}$$

The CRLB for the estimate of the mth metabolite is then:

$$CRLB_m = \sqrt{(F^{-1})_{m,m}} \tag{2.7}$$

In equation 2.6, the matrix B^{total} refers to the m by r matrix containing the simulated metabolite signals where r is the size of the spectral ROI. Each column of B^{total} is a vectorized from of the 2D spectrum for a given signal. σ^2 denotes the standard deviation of the real part of the noise. The assumptions implied in equation 2.7 are that the noise is white, the fitting model is correct and an optimal fit is obtained[14].

An interesting result is that the CRLB is dependant on the estimated noise and orthogonality of the basis set. This in part explains why quantification of 2D spectroscopy can achieve lower bounds than 1D spectroscopy as the introduction of a second dimension improves the orthogonality of the columns of the matrix B^{total}. However; caution should be taken in comparing 1D with 2D quantification accuracy as the number of scans in a 2D experiment always many more than an 1D experiment [7].

Chapter 3

Compressed Sensing Theory and Application to 2D NMR

It is difficult to overstate how the development of signal theory over the past century has contributed tremendously to the technological achievements of modern civilization. This is evidenced by its ubiquitous application to nearly every human endeavour including medicine, automation, communication, industrial processing and travel by land, air and sea. The great power of signal theory lies in a framework that allows information to be processed, analyzed and transformed, allowing natural and synthetic phenomena to be exploited and utilized in ways that could not be realized in their raw form.

3.1 Harmonic Analysis

Consider the periodic, continuous-time signal, $f(t)$. Without loss of generality, $f(t)$ is treated as a complex-valued function of the time variable, t. A signal is defined as

being periodic if for T being a fixed, positive real number:

$$f(t+T) = f(t) \qquad for\ all\ t, \quad -\infty < t < \infty \tag{3.1}$$

Note that equation 3.1 holds for also for nT where n is any positive integer. The smallest positive number, T, for which (3.1) holds is denoted the fundamental period.

Owing to the pioneering work of the French physicist, Jean Baptiste Fourier (1768-1830), $f(t)$ can be expressed as an infinite summation of complex exponential functions such that the Fourier Series of $f(t)$ is given by:

$$f(t) = \sum_{k=-\infty}^{\infty} c_k e^{jk\omega_0 t}, \quad -\infty < t < \infty \tag{3.2}$$

where

$$c_k = \frac{1}{T} \int_0^T f(t) e^{-jk\omega_0 t} dt, \quad k = 0, \pm 1, \pm 2, ... \tag{3.3}$$

In the formulation above, $\omega_0 = \frac{2\pi}{T}$ is denoted the fundamental frequency. As stated previously, the Fourier Series is only defined for periodic functions but only if they satisfy the Dirichlet conditions: $f(t)$ is absolutely integrable and has only a finite number of maxima, minima and discontinuities over any period. Virtually all interesting signals that arise in engineering satisfy these conditions, but very few are truly periodic. The generalization of the Fourier Series to aperiodic signals is given by the inverse Fourier transform:

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega) e^{j\omega t} dt, \quad -\infty < t < \infty \tag{3.4}$$

where the Fourier Transform of $f(t)$, $F(\omega)$, is given by:

$$F(\omega) = \int_{-\infty}^{\infty} f(t)e^{-j\omega t}dt, \quad -\infty < \omega < \infty \tag{3.5}$$

The signals $f(t)$ and $F(\omega)$ constitute a Fourier Transform pair. It should be noted that the Fourier Transform is a linear operation that can be reversed without loss of information. Fourier pairs are denoted in the form:

$$f(t) \rightleftharpoons F(\omega) \tag{3.6}$$

The Fourier Transform has numerous uses including spectral analysis, solving ordinary differential equations and computing the impulse response of a system. These owe to the various unique properties of the Fourier Transform, some of which will be explored further in the following sections.

3.2 Nyquist Sampling Theorem

In reality, signals of interest cannot be easily described in a compact, continuous mathematical form, $f(t)$. The digital revolution was built upon the notion that by measuring the signal for a finite number of regularly spaced time points, then the information content of the underlying signal is preserved. This process is known as sampling and virtually all types of media are captured in this way. The question of how often a signal must be sampled is addressed by the famous Nyquist Sampling Theorem which states that the sampling frequency must be at least twice the largest frequency component of a band-limited signal. This claim can be demonstrated by

introducing the sampling operator and observing its effect on the frequency-domain representation of a signal.

Idealized sampling ignores the intricacies of the sampling process and can be formalized mathematically by multiplication with a sampling function or Dirac train. The Dirac delta function is defined as:

$$\delta(x) = \begin{cases} \infty, & x = 0 \\ 0, & x \neq 0 \end{cases} \qquad \int_{-\infty}^{\infty} \delta(x)dx = 1 \tag{3.7}$$

For a sampling period, T_s, the sampling function is given as:

$$\Delta(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT_s) \tag{3.8}$$

The sampling function can be represeneted as a Fourier Transform Pair:

$$\sum_{n=-\infty}^{\infty} \delta(t - nT_s) \rightleftharpoons \frac{2\pi}{T_s} \sum_{k=-\infty}^{\infty} \delta(\omega - \frac{2\pi k}{T_s}) \tag{3.9}$$

In words, the Fourier Transform of a train of impulse functions separated in time by T_s is a train of impulse functions separated in angular frequency by $\omega_s = \frac{2\pi}{T_s}$. An important property of the Fourier Transform that has tremendous consequences for sampling is the multiplication-convolution duality, such that:

$$F\{f(t) \cdot g(t)\} = F\{f(t)\} * F\{g(t)\} \equiv F(\omega) * G(\omega) \tag{3.10}$$

The sampling operation is formalized mathematically as time domain multiplication with the sampling comb. In the frequency-domain, this can be represented as a

convolution of the signal spectrum with the Fourier Transform of the sampling operator. The result in continuous-time analysis is an infinite number of replications of the frequency-domain signal, separated by the sampling frequency. Nyquist theorem claims that in order to recover a sampled signal, the sampling frequency must be at least twice as high as the largest frequency component of a band-limited signal. This prevents frequency-domain overlapping of the sampled signal from which the original signal can be recovered through low-pass filtering the base-band replicate. If the sampling frequency is lower than this condition, then the result is known as aliasing, a situation in which the signal is distorted by the interference of higher and lower frequency components.

3.3 The Discrete Fourier Transform

The Fourier Transform of a sampled signal can be expressed as a Discrete-Time Fourier Transform (DTFT). Because the continuous-time signal can be represented as a sequence, this reduces the integral of equation 3.3 to a summation, giving the DTFT as:

$$F(\omega) = \sum_{n=-\infty}^{\infty} f[n]e^{-j\omega n} \tag{3.11}$$

In a real-world scenario, a signal is sampled and digitized, leaving only a finite number of measurements. This changes the limits of the summation in 3.11 to only include the samples that fall within the sampling window. Note that in the DTFT operation, the frequency-domain representation is a continuous function in frequency. For this reason, it is useful then to introduce an operator for mapping a finite number of measurements to a finite number of frequencies. This operation is known as the

Discrete Fourier Transform or DFT.

One interpretation of the DFT is stems from Fourier Series analysis. Recall that the Fourier Series is valid only for periodic functions. By treating a finite sequence of measurements as being a single period of a periodic function, then the DFT can be viewed as Fourier Series analysis for discrete-time signals. The notion of sampling frequency is transparent to the DFT operation, and the fundamental period is taken as the total number of samples, N. The DFT can be represented as:

$$F[k] = \sum_{n=0}^{N-1} f[n] e^{-j\frac{2\pi}{N}nk} \quad k = [0, N-1] \tag{3.12}$$

Before discussing the inverse operation, it is useful to introduce the DFT in the compact matrix-vector multiplication form. For a given discrete frequency, k, the associated coefficient, $F[k]$, is computed as a dot product of the signal vector, $f[n]$, with the complex exponential vector. For all N discrete frequencies, the DFT operation can be summarized in the following form:

$$\begin{pmatrix} F[0] \\ F[1] \\ F[2] \\ F[3] \\ \vdots \\ F[N-1] \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & W & W^2 & \cdots & W^{N-1} \\ 1 & W^2 & W^4 & \cdots & W^{N-2} \\ 1 & W^3 & W^6 & \cdots & W^{N-3} \\ \vdots & & & & \\ 1 & W^{N-1} & W^{N-2} & \cdots & W \end{pmatrix} \begin{pmatrix} f[0] \\ f[1] \\ f[2] \\ f[3] \\ \vdots \\ f[N-1] \end{pmatrix} \tag{3.13}$$

where $W = e^{-j\frac{2\pi}{N}}$. Let W represent the N-point discrete Fourier Transform matrix. Due to complex symmetry, the inverse transform can be shown through the following

identity:

$$W^{\dagger}W = WW^{\dagger} = N \cdot I \tag{3.14}$$

The discrete Fourier Transform is made unitary by scaling the DFT matrix by $\sqrt{\frac{1}{N}}$. For following discussions, this scaling is assumed. The DFT matrix can be viewed as an orthonormal basis, and the discrete Fourier Transform as an orthonormal or unitary transform. This necessitates that through forward/reverse transformation, the signal is preserved. Each element of the solution vector is given by the dot product of the signal with a row vector in the DFT matrix. Stated in another way, the nth discrete Fourier coefficient is the fractional content of the nth DFT column or basis vector in the signal.

3.4 Compressed Sensing

Consider the classical linear measurement problem, where the measurements $y = Ax$. If vectors y and x complex vectors of dimension n, then A is an n by n matrix that maps the process x to the measurements y. Assuming there is no noise, then the signal x is recovered uniquely through the linear measurements Ax such that $x = Inv(A)y = Inv(A)Ax = Ix = x$ where I is the identity matrix and $Inv(\cdot)$ is the matrix inverse operator. This procedure can be viewed as an extension of Nyquist theorem such that to reconstruct an n-point signal, n measurements are required.

In the situation where y is a vector of dimension m, where y, then the system of equations $y = Ax$, where A is an m by n matrix, is an underdetermined system and it cannot be solved in general. Compressed Sensing (CS) is a mathematical framework that claims that if there exists some specific prior knowledge on the structure of x,

then the underdetermined system of equations has a unique solution [17].

In 2004, Emmanuel Candes, Justin Romberg and Terence Tao (CRT) published the first guarantees for uniquely reconstructing an n dimensional complex signal from a set $m < n$ randomly chosen frequency samples. The authors considered the problem of randomly undersampling in frequency whereas in the application of interest to this thesis, the undersampling is performed in time. Since the DFT is a unitary transform, all of the claims made by CRT directly apply.

Consider the complex vector, x, of dimension n. Ignoring noise, the measurement process can be cast as linear measurements of the form $y = Wx$ where W is the inverse Fourier Transform matrix. Conversely, the digitized complex sequence y can be Fourier Transformed to reveal the signal spectrum, x. CRT showed for prime n, as long as the number of measurements of signal y is greater than or equal to twice the number of non-zero frequency components of the solution vector x, then the measured signal is unique. For example, if vector x contained two frequency components with the rest zero, then it would be described as being k-sparse with k equal to two. CRT claim in [17] that x is uniquely captured as long as the number of measurements, m, is at least $2k$, or in this example, four. This means that if two separate signals were sampled in the same manner and the samples were identical, then the underlying signals must also be identical. Uniqueness is a central property that is necessary to hope to solve such an underdetermined system.

Since a k-sparse signal is uniquely captured from at least $2k$ measurements, where the matrix A is the partial inverse discrete Fourier Transform, then this suggests that the underlying signal can be uniquely reconstructed according to:

$$min \parallel x \parallel_{l_0} \qquad s.t.\ y = Ax \tag{3.15}$$

The cost function references the zero-norm, which is in fact a pseudo-norm and is more accurately classified as a metric. The zero-norm metric returns the sparsity of the signal, counting the number of non-zero elements. In words, the cost function in (3.15) states that the unique solution to the underdetermined system of equations, $y = Ax$, can be solved by finding the vector x which has the fewest non-zero components and agrees with the measured values. In practice, this problem is computationally intractable and has no efficient solution. The problem can be solved by exhaustively searching through every combination of the signal support. For each possible combination of the signal support for each possible level of sparsity, up to half the number of measurements, the associated submatrix of the measurement matrix is inverted and the solution is checked against the measured values. For even modest signal sizes, this procedure is has exponential complexity and no robustness to noise[17].

The claims cited in [17] are only valid for prime n since for arbitrary n, there exists the possibility of the existance of signals that cannot be uniquely captured. An illustrative counter example is the Dirac comb of length n where n is a perfect square, and the impulses are separated by the square root of n. This particular signal is invariant through the unitary discrete Fourier Transform, such that there exists the possibility that all of the measurements may be taken from time points where the signal is zero. By assuming that the signal support and measurements are chosen at random, then sparse signals can be reconstructed with probability approaching one[17]. CRT showed that by relaxing the l_0 norm to the l_1, that a randomly undersampled sparse signal can uniquely reconstructed according to:

$$min \parallel x \parallel_{l_1} \qquad s.t.\ y = Ax \tag{3.16}$$

The problem of (3.16) is now commonly known as basis pursuit [20], and can be efficiently solved in the complex case as a second order cones problem. The cost of relaxing the combinatorial problem in (3.15) to the convex problem in (3.16) is an increased number of measurements, proportional to the sparsity of the solution vector, k, and the introduction of a logarithmic dependance on the signal dimension, n. It was proven originally in [18] that the number of m measurements required to reconstruct an n-dimensional k-sparse vector with overwhelming probability is upper bounded such that:

$$m = O(k \cdot log^6 n) \tag{3.17}$$

The logarithmic exponent was improved in [19] by reducing it to $m = O(k \cdot log^4 n)$. Although the estimate derived in [19] is currently the best known bound, it is conjectured that the optimal bound is $m = O(k \cdot log n)$ as is the case for considering a fixed x and all possible combinations of undersampling x, i.e. The non-universal case.

It is important to note that the fundamental claims of compressed sensing are not deterministic but probabilistic in nature. The expanding mathematical framework answers the question for a k-sparse signal with randomly distributed support and randomly indexed measurements, what is the *probability* of recovering the original signal? Certainly in reality, when the concern is to faithfully measure a signal, it is important to operate in the region where this recovery probability very nearly approaches one.

3.5 The Random Fourier Measurement Matrix

In order to formalize the definition of the measurement matrix, A, let Ψ denote the n by n dictionary matrix, and let Φ denote the m by n sensing matrix. In these terms, A is the m by n measurement matrix such that $A = \Phi\Psi$. In the context of partial Fourier measurements, the sensing matrix Φ, is constructed as any set of size m of the rows of the n by n identity matrix. Consider the following figures as a concrete illustration: Figure 3.1 depicts an interpretation of fully sampling a time-domain

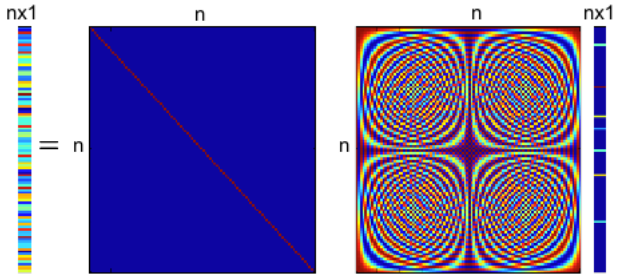

Figure 3.1: Complete Time-Domain Sampling

signal. Since Ψ is a unitary matrix, sampling the time-domain signal can be viewed as being equivalent to measuring the orthonormal projection of the frequency-domain signal as an n-point set of time-domain coefficients.

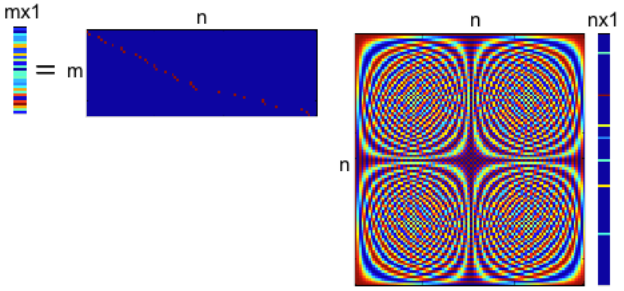

Figure 3.2: Time-Domain Undersampling

Figure 3.2 illustrates the process of temporal undersampling. The matrix-vector

product of the n by n inverse DFT matrix with the n by 1 sparse vector x is equivalent to the fully sampled time-domain signal as in Figure 3.1. Undersampling is performed by selecting m of the n entries of the time-domain vector y such that $y = \Phi y_{full} = \Phi\Psi x$. The measurements, y, in figure 3.3 are equivalent to those depicted in figure 3.2,

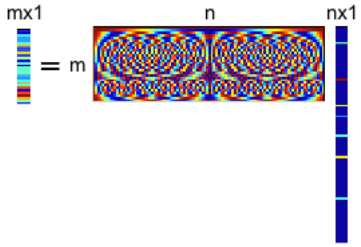

Figure 3.3: Compressed Sensing

but for completeness, the m by n measurements matrix $A = \Phi\Psi$ is shown. Assuming that the system can be solved, the framework of compressed sensing asserts that the measurements represent a form of compact representation of the underlying signal. This measurement scheme is analogous to image compression, as a high dimensional signal can be accurately represented as a collection coefficients in a lower dimensional space. The paradigm of compressed sensing is that it is possible to exploit this property directly, and acquire compressed measurements.

3.6 UUP, Compressibility and Noise

The uniqueness theorem for random Fourier sampling of k-sparse signals asserted that solving (3.15) may be performed as a combinatorial search, inverting each m by d minor of the partial Fourier matrix, A_d,for each possible set of columns d pertaining to the support of the solution vector. The Uniform Uncertainty Principle (UUP)

states that the following inequality:

$$\frac{1}{2}\frac{m}{n} \parallel x \parallel_{l_2}^2 \leq \parallel Ax \parallel_{l_2}^2 \leq \frac{3}{2}\frac{m}{n} \parallel x \parallel_{l_2}^2 \tag{3.18}$$

holds for *any* k-sparse signal x, with k obeying the bound in (3.17), where m is the number of measurements, and n is the signal dimension. This powerful result can be interpreted such that as long as any k-sparse signal is sufficiently sampled, then the energy of its m measurements is at most only a fraction of the energy of the underlying n dimensional signal. In other words, with overwhelming probability, the sparse spectrum x cannot be concentrated in time. This is the uncertainty principle, such that any of the possible measurements are equally (un)important. The uniform part of the UUP speaks to the strong claim that (3.18) holds for *every* possible signal support of reasonable size. Inequality (3.18) is equivalent to its original formulation, which is a statement about the bounds of the minimum and maximum eigenvalues of the square matrix formed by $A_d^* A_d$ [18, 20].

Although the UUP provides an elegant grounds for why l_1 minimization works for recovery of k-sparse signals from incoherent measurements, it does not account for realistic signal models. In the absence of noise, natural signals when projected into a sparsifying basis usually consist of a small subset of dominant coefficients with the rest of negligible value. Lossy compression techniques such as JPEG-2000 and MP3 manage to retain signals that are nearly indistinguishable from their originals by exploiting the asymptotic sparsity of natural signals following a sparsifying transform, and adaptively truncating the smallest coefficients.

In [18], the authors showed that the error of recovering an approximately sparse

signal whose coefficients decay according to a power law is on comparable with knowing everything about the signal and setting all but the k-largest terms to zero. This result was extended further in [20] by considering the robustness of l_1 minimization to measurement noise. The authors showed that for a bounded, unknown perturbation, the recovery is stable with a reconstruction error proportional to the noise level. The noise bound is incorporated into the convex minimization problem solving:

$$min \parallel x \parallel_{l_1} \qquad s.t. \parallel Ax - y \parallel_{l_2} \leq \epsilon \tag{3.19}$$

Problem (3.19) is referred to as basis pursuit denoising and it seeks the sparsest solution that satisfies the equation in a minimum euclidian distance within a tolerance. The tolerance, denoted ϵ, is termed the noise level such that if $y = Ax + e$, ϵ is taken as the l_2 norm of the unknown perturbation vector.

The error for compressed sensing of noisy, compressible signals is bounded by:

$$\parallel x_0 - x^* \parallel_{l_2} = O(\sqrt{\frac{n}{m}} \cdot \epsilon + \frac{\parallel x_0 - x_{0,k} \parallel_{l_1}}{\sqrt{k}}) \tag{3.20}$$

This result can be read that the reconstruction is bounded by the noise in the measurements, as well as the best-k-term approximation error in the signal, denoted $x_{0,k}$[20].

3.7 Algorithms that solve y = Ax

Basis pursuit (3.16) and the denoising variant (3.19) can be solved by a linear program in the real valued case or a second order cones problem in the complex-valued case.

Modern interior-point solvers can be exploited to solve the problem more efficiently than general purpose convex optimization [17]. Iterative greedy algorithms are a family of approaches that seek to minimize the energy of the residual in (3.16) subject to a sparsity enforcing constraint. Refer to [21] for an overview of these approaches. Convex relaxation of the l_0 norm to the l_1 norm has been shown to provide the sparsest solution [22] and hence, l_1 can be viewed as a proxy for sparsity. Algorithms which attempt to minimize the l_p norm for $0 < p < 1$ sacrifice the convexity of problem (3.16) but can be solved more efficiently than the combinatorial search of l_0 minimization. The authors in [23] show that l_p minimization improves the theoretical guarantees with respect to noise stability and compressibility beyond l_1 minimization.

One of the most straight forward implementations of l_p minimization is the Iteratively Re-Weighted Least Squares (IRLS) family of algorithms. To review the least squares problem, the minimum l_2 solution to $y = Ax$ is given by:

$$x^* = A^\dagger y \tag{3.21}$$

where $A^\dagger$ denotes the Moore-Penrose pseudoinverse of A. Since, the rows of A are orthonormal, $A^\dagger$ can be computed as:

$$A^\dagger = A^T(AA^T)^{-1} \tag{3.22}$$

Note that the quantity within the brackets is the n by n identity matrix when A is the m by n partial DFT matrix. The solution in (3.21) is equivalent to the least squares solution of $y = Ax$. The Moore-Penrose psuedoinverse solution is the projection of y onto the row space of A. The minimum l_2 solution seeks the solution that minimizes

the energy of the solution vector. The result is that the missing measurements are set to zero. This is the least committal solution, but there is no reasonable assumption as to why the missing measurements should be zero.IRLS makes use of that fact that for any $p,\ 0 < p < \infty$:

$$\| x \|_p^p = \sum_{n=1}^{N} \frac{| x_n |^2}{| x_n |^{2-p}} \tag{3.23}$$

and that weighted least squares problem,

$$min\{\sum_{n=1}^{N} d_n x_n^2\} \quad s.t.\ y = Ax \tag{3.24}$$

has the closed form solution:

$$x = D^{-1}A^T(AD^{-1}A^T)^{-1}y \tag{3.25}$$

where D is an n by n matrix whose diagonal entries are given by:

$$d_n = \frac{1}{| x_n |^{2-p} + \epsilon} \tag{3.26}$$

The algorithm proceeds iteratively where the weights are updated from the previous iteration until a set number of iterations. ϵ here is a small constant used to ensure stability of the problem by avoiding division by zero. The differences between various implementations of IRLS are the manner in which the weights are updated, the l_p norms used and the possibility of updating the stability parameter. The objective function can also be modified to include a fidelity parameter that incorporates a tradeoff between data consistency and sparsity [24].

Since the first works on compressed sensing were published in 2004, the field has

grown at a rapid rate, producing many variants of algorithms that solve the system $y = Ax$ when x is sparse. As alluded to in [21], sifting through the dense forest of algorithmic papers is a tall order for engineers and applied scientists who hope to apply compressed sensing in real-world scenarios. Given the qualitative nature of this thesis, a thorough review of the numerous implementations of l_1 and l_p minimization schemes are omitted. The primary concern here is that the algorithm, whatever it may be, must allow the use of complex-valued inputs and it must converge to a reasonable solution within an acceptable amount of computation time. General purpose convex solvers (for l_1 minimization) and simple implementations of IRLS (for l_p minimization) both satisfy these loose requirements.

An additional formulation of the sparsity seeking problem which satisfies the requirements of this work is known as the smoothed l_0 algorithm (SL0). SL0 attempts to solve the l_0 minimization problem in (3.15) by approximating the l_0 with a smooth gaussian function with parameter σ. Compared to general convex solvers, SL0 is extremely fast, making it suitable for large scale problems. SL0 is one particular example of a specialized algorithm that allows the use of functional definitions of the measurement process in its implementation, meaning that the Fast Fourier Transfrom (FFT) can be employed to vastly scale down computation time. For a detailed description of SL0 refer to [25].

3.8 Application of Compressed Sensing to 2D NMR

The modern application of NMR is primarily focused on multidimensional experiments due to the richness of information which may be extracted from such experiments. As JPRESS can be used in order to reduce the spectral clutter by increasing

the dimensionality of the data set (section 2.2), two, three, four and even higher dimensional experiments are routinely performed on protein NMR experiments in order to elucidate the composition of such complex species comprised of thousands of spins and their local interactions. In the liquid state, the acquisition of a single transient along with an appropriate recovery delay is on the order of 1-5 seconds. Since the NMR signal is short lived, there would be no improvement in measurement time by undersampling the direct dimension and it would require specialized hardware. In this sense, these measurements come for free. However, the total experiment time in a two dimensional NMR experiment,for example, with repetition time denoted TR is given by:

$$T_{total} \propto TR \cdot NEX \cdot N_1 \tag{3.27}$$

where the NEX refers to the total number of signal averages and N_1 denotes the resolution of the indirect dimension. The desired indirect resolution dominates the total experiment time and is generally kept within the range of 64-128 for *in vivo* studies simply to limit the total scan duration. Although there is less pressure to keep the experiment time down for conventional NMR studies, the total time scales exponentially with data dimensionality meaning experiments can potentially last hours to days in higher dimensional experiments. The long duration of multidimensional NMR has motivated the community to improve throughput by reducing acquisition time [26, 27, 28].

The first published work demonstrating the application of compressed sensing to 2D NMR was in 2007 [26], proposing to solve the basis pursuit denoising problem (3.19) in the case where the sensing matrix was the partial DFT matrix and the dictionary matrix was an orthogonal discrete wavelet transform. The author presented

results for undersampling a noiseless synthetic data set by a factor of two, claiming consistent results of an RMSE of 0.5 for various resolutions and numbers of peaks. Additionally, the author demonstrated a linear relationship for reconstruction error versus the number of peaks. This result is in good agreement with the bounds in (3.17).

In 2011, Qu *et al.* showed that using the wavelet transform as a sparsifying transform for 2D spectroscopy is actually suboptimal in comparison to the canonical identity basis. The main result was that most two dimensional data sets are self sparse and don't require an additional sparsifying transform. The authors demonstrated this claim by showing that the the sorted coefficients in the identity basis decay faster than in an orthogonal wavelet basis. This result is in line with the guarantees for compressed sensing concerning approximate sparsity. Additionally, the notion of using a wavelet transform as a sparsifying basis is misguided as the orthogonal wavelet transform was designed for compactly representing piece-wise smooth signals [29], a signal model that does not apply well to self sparse spectra.

Holland *et al.* published the first results of applying compressed sensing to 3D NMR of biological proteins in 2011 [30]. The focus of the paper was on the consistency of amino acid back bone assignment following reconstruction from only a fraction of the full data. In cases where less than 30

In 2012, Hyberts *et al.* published results that applied l_1 minimization to multidimensional NMR [31]. The authors demonstrated passable reconstructions of 3D and 4D datasets by sampling less than 1% of the nyquist grid. Kazimierczuk *et al.* offered the first actual attempt at constructing a performance measure of the l_1/l_p minimization approach by comparing the reconstruction quality of a set of synthetic

signals for a large number of different sampling schedules [32]. This work explicitly acknowledges the fact that for the total number of sampling points held constant, the quality of the reconstructions over a large number of permutations is in fact variable. This observation stems from the probabilistic nature of compressed sensing. In order to assess the reliability of l_1/l_p minimization, the various reconstructions were classified in a binary way such that if the l_2 norm of the residual was less than an arbitrary fraction of the true l_2 norm of the signal then the reconstruction was deemed perfect. The authors showed that beyond a data set dependant number of samples, the probability of perfect reconstructions saturated at 100%. Additionally, the authors compared various data sets with varying peak crowdedness and showed that this saturation point was proportional to sparsity. Finally, they showed that in most cases that l_p minimization outperformed l_1 minimization with respect to the minimum sampling requirement as well as computation time.

Chapter 4

Evaluation of Compressed Sensing Reconstruction of Simulated JPRESS Data

In order to investigate the feasibility of applying Compressed Sensing to the acquisition of JPRESS data, model spectra was generated reflecting various brain metabolites. The following sections attempt to address several important considerations such as the limit of undersampling, the choice of sampling schedule and the robustness to noise.

4.1 Materials

The simulated JPRESS pulse sequence was implemented using the Spinach[12] library for Matlab (MATLAB Release 2012a, The MathWorks, Inc., Natick, Massachusetts, United States). Sampling periods and pulse timings were designed to reflect the

JPRESS implementation on the GE 3T Signa platform (General Electric Medical Systems, Miluake, U). Refer to section 5.1 for the specific acquisition parameters. Synthetic metabolites were designed according to the contents of the GE Spectroscopy phantom which is depicted in table 4.1.

Symbol	Chemical Name	Concentration
NAA	N-acetyl-L-apartic acid	12.5mM
Cr	Creatine hydrate	10mM
Ch	Choline chloride	3.0mM
mI	Myo-inositol	7.5mM
Glu	L-glutamic acid	12.5mM
Lac	DL-lactic acid	5mM

Table 4.1: GE Spectroscopy Phantom Contents

Chemical shifts and scalar coupling values were chosen according to reported values [8]. Water was also included in the synthetic data set to simulate the effect of imperfect solvent suppression. Figure 4.1 depicts an overlay of the 1D spectra for each metabolite including water.

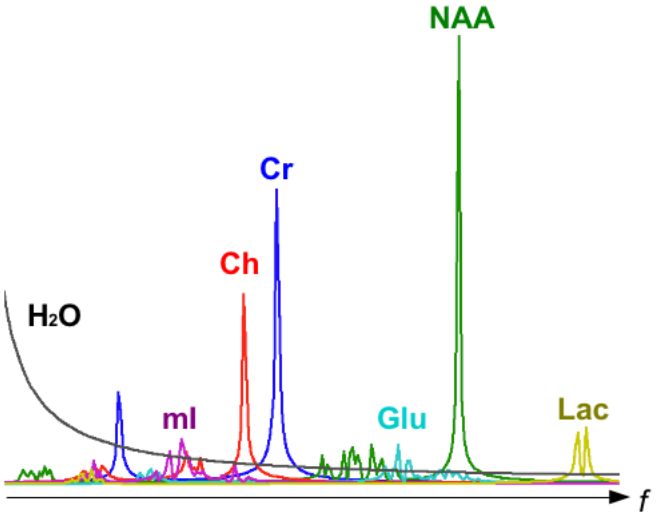

Figure 4.1: Simulated Metabolite Spectra (1D)

4.2 Comparison of Sparsity Seeking Algorithms

The simulated time-domain metabolite data was subsampled uniformly random distribution and then Fourier Transformed in the direct dimension. The problem was divided into a sequence of $n_2 = 1024$ reconstructions whose solution vectors are each of dimension $n_1 = 100$. n_2 is the number of points in the direct dimension and n_1 is the number of points in the indirect dimension in the fully sampled data set. Three sparsity seeking algorithms were compared on the basis of average run time as well as Reconstruction SNR (RSNR) defined as:

$$RSNR = 20 \cdot log_{10} \frac{\| x_0 \|_{l_2}}{\| x_0 - x^* \|_{l_2}} \tag{4.1}$$

where x_0 is the true solution and x^* is the recovered signal. The algorithms considered were the direct implementation of basis pursuit using the constrained optimization library CVX[**?**], an implementation of IRLS that approximates the zero-norm [32] and SL0[36] which was modified the make use of the Fast Fourier Transform. The results for a single trial are plotted below in figure 4.2. A trial was defined as a single reconstruction using set of samples for a given sampling rate. In the figures that follow, the sampling rate is defined as the dimensionless fraction of the number of measured points compared to the total signal dimension. It is impossible to make any strong conclusions as to which is the best algorithm when considering the RSNR for an isolated trial. Only one sampling schedule was selected from a uniform random distribution for each sampling rate. Due to the probabilistic nature of Compressed Sensing, the choice of samples is expected to impact the reconstruction unless the signal truly sparse and is sufficiently oversampled. Observe in figure 4.2 how the

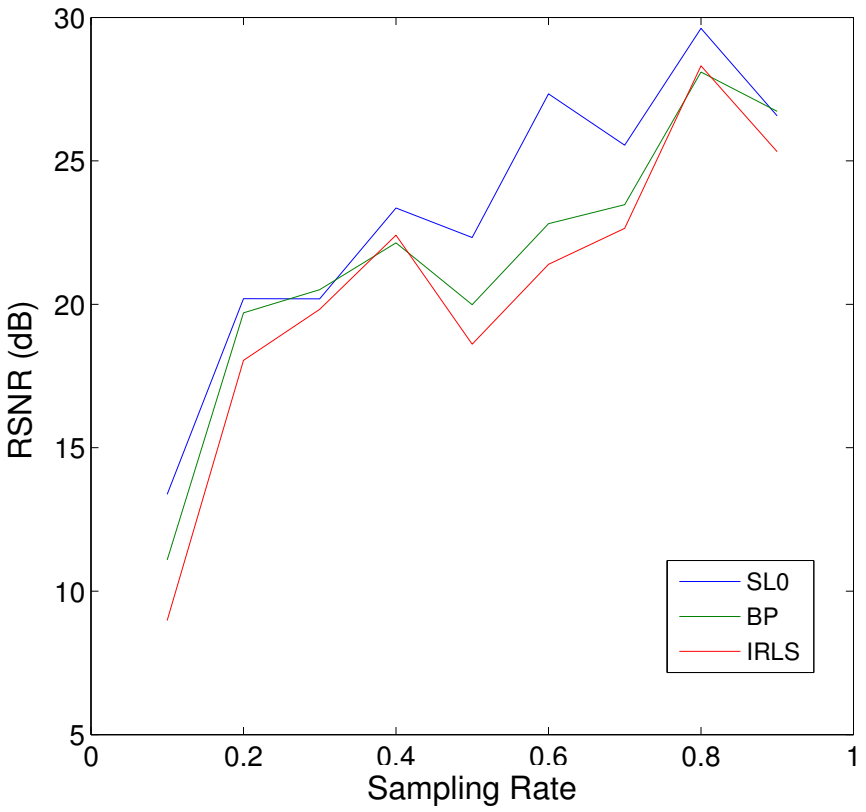

Figure 4.2: RSNR vs Sampling Rate

RSNR is higher for all algorithms when using 80% of the samples rather than 90%. This shouldn't be interpreted to conclude that sparsity promoting algorithms are more efficient with fewer samples, rather it might imply that the random samples in that instance were better than the random samples for the higher sampling rate trial. The important observation is that all algorithms performed better in that instance. This suggests that despite the diversity in the algorithmic approaches, the outputs for each implementation are relatively consistent given the same input.

Figure 4.3 depicts the running time taken for each reconstruction with respect to the sampling rate. In all cases, the running time increases with the size of the

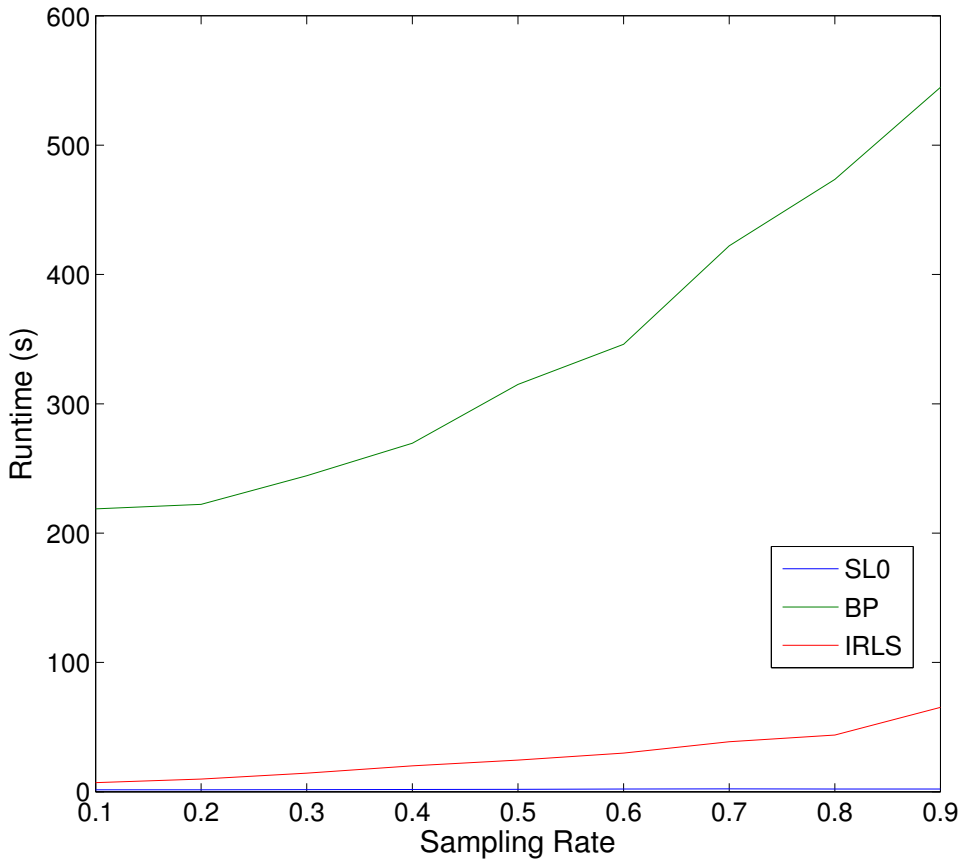

Figure 4.3: Runtime vs Sampling Rate

problem. SL0 is the clear winner, followed by IRLS with the CVX implementation of BP being the slowest. A full complexity analysis is beyond the scope of this thesis; however, some general claims can be made. The implementation of IRLS used here proceeds for a predefined number of iterations, whereas BP ans SL0 both have stopping criteria. For a given number of iterations and problem size, the runtime of IRLS will always be the same. In this way, IRLS can be sped up by reducing the number of iterations while potentially sacrificing accuracy. 50 iterations were chosen although 10-20 would have sufficed. This was to ensure a fair accuracy comparison

with the other two algorithms. The minimum and maximum runtimes for IRLS, corresponding with the smallest and largest system of equations, was 7 seconds and 65 seconds respectively.

The CVX implementation of BP was the slowest by a large margin. Specialized sparsity seeking algorithms were conceived on the basis that generalized solvers are too slow for even modest problem sizes. The minimum and maximum runtimes for BP in CVX was 218 seconds and 544 seconds respectively. SL0 was the only specialized algorithm considered, developed to quickly find accurate sparse solutions. The ability for SL0 to make use of the Fast Fourier Transform contributes greatly to its speed in this analysis. The minimum and maximum runtimes for SL0 was 1.4 seconds and 2.1 seconds respectively.

The focus of this thesis is to explore the feasibility of applying Compressed Sensing to the acquisition of JPRESS data. For the following analysis, SL0 is taken as the algorithm of choice. This choice was made on the basis of runtime alone in which SL0 was the clear winner. To provide a thorough treatment of the considerations that follow, a variety of repetitive problems must be solved. It is important to decouple the concept of seeking sparse solutions and the actual implementation. As mentioned, the focus is evaluating the application of Compressed Sensing to JPRESS rather than determining the best algorithm for application of Compressed Sensing to JPRESS.

4.3 Evaluation of Zero-Padding

The coefficients of the discrete Fourier Transform (DFT) can be interpreted as being the result of convolving the discrete time Fourier Transform (DTFT) with a Dirichlet kernel and then sampling. This means that unless the time domain sinusoids are

integral frequencies, spectral leakage will occur that follows the Dirichlet envelope. Because of this, the DFT of complex sinusoids results in signals that aren't truly sparse but display compressibility. Decaying exponentials further compound the issue of sparsity owing to the Lorentzian lineshape. These issues can be remedied to an extent by zero-padding the signal, which provides samples from the DTFT at a higher rate. The sensing matrix in this case is no longer the DFT, but can be considered as a redundant frame of complex sinusoids [33]. Figure 4.4 depicts the relationship between RSNR and the amount of zero-padding. This experiment considers a uniform random undersampling with sampling rate 0.5 and the extent of zero-padding is taken in integer multiples of the direct and indirect dimension size. This experiment, as well as all that follow, use a direct 2D implementation of SL0 which make use of the two dimensional Fast Fourier Transform. During each iteration of the SL0 algorithm, the solution vector is reshaped, transformed, masked with zeros at the missing samples and then vectorized into a single $n_1 \times n_2$ vector.

Figure 4.4a presents a qualitative description of the relationship between RSNR and zero-padding. The baseline RSNR is 19.38 dB. Figure 4.4b and c depict the gradient of figure 4.4a along the indirect and direct dimension respectively. Caution must be used when interpreting the gradient images alone since the colour scales are identical in all images, but the range and units are not. Figure 4.4d depicts the runtime of the experiment with respect to the extent of zero-padding. Figure 4.4c clearly shows strong gains in RSNR by zero-padding the indirect dimension by a factor of one, with marginal gains when zero-padding by a factor of two. The trend in figure 4.4b is less obvious as the gains in RSNR don't seem to apply until there is some zero-padding in the indirect dimension. The gradients in both images show

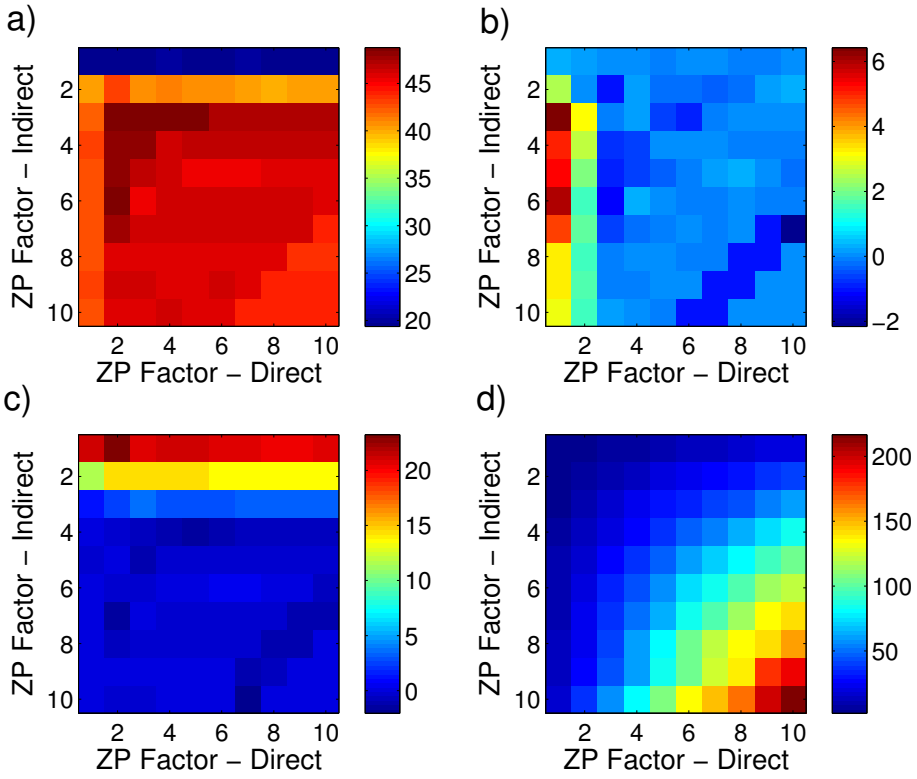

Figure 4.4: a) RSNR as a function of the extent of zero-padding in each dimension b) The gradient of a) with respect to the direct dimension c) The gradient of a) with respect to the indirect dimension d) Runtime as a function of the extent of zero-padding in each dimension

not only diminishing marginal returns to zero-padding, but even diminishing returns. The peak RSNR is obtained when the amount of zero-padding in each dimension is a factor of three, corresponding to 48.77 dB. The RSNR when both dimensions are zero-padded by a factor of 10 is 43.89 dB. Regardless of the amount of zero-padding, it never reduced the quality of the reconstruction compared to the baseline. Figure 4.5a plots the diagonal entries of figure 4.4a, corresponding to linearly increasing the amount of zero-padding in each dimension. The runtime in seconds is plotted against the zero-padding factor in figure 4.5b.

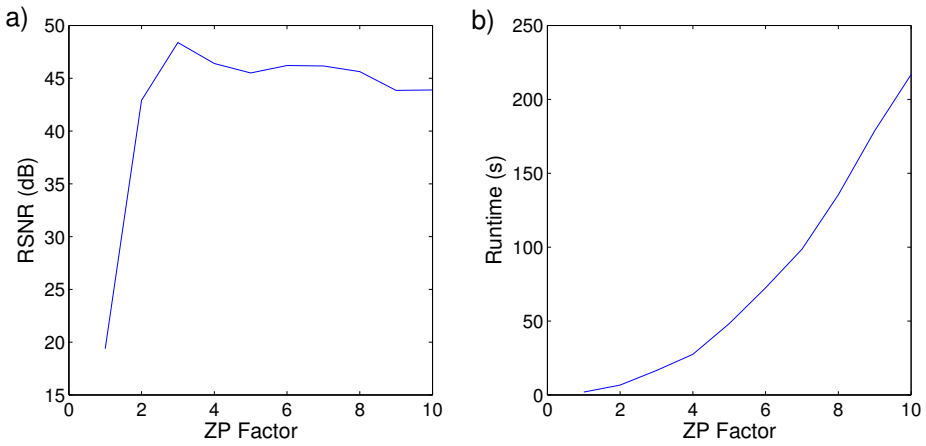

Figure 4.5: a) RSNR vs ZP Factor b)Runtime vs ZP Factor

4.4 Evaluation of Sampling Schedules

The quality of Compressed Sensing reconstruction of undersampled 2D spectroscopy is strongly dependant on which particular samples are chosen. This statement is easily justified by observing the variation in reconstruction quality while the total number of samples is fixed. To this end, there have been several investigations into the impact of sampling into non-linear reconstructions [34]. In this section, two sampling schemes are investigated; uniform random sampling and sinusoidally modulated Poisson Gap (PG) sampling. PG sampling was determined by Hyberts *et al.* to offer the highest average case reconstruction quality and is based off the groups observations in random undersampling. They concluded that the most poorly reconstructed time points were in regions devoid of samples and that the reconstruction overall is poorest when there are large gaps at the beginning of the schedule. The PG method attempts to randomly place gaps in the sampling schedule, maximizing the density of early samples and maintaining relatively small gaps across the entire schedule. Refer to [34] for details. Typical sampling schedules for each approach are depicted in figure

4.6.

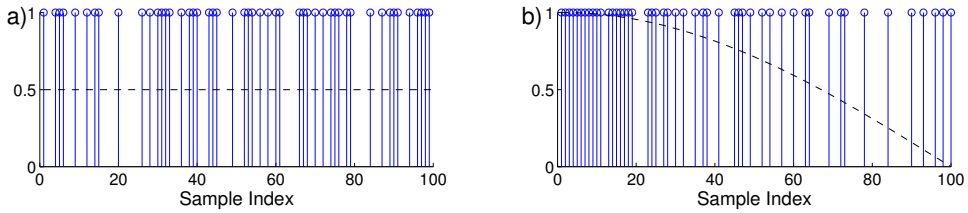

Figure 4.6: a) Uniform Random b)Sinusoidally Modulated Poisson-Gap

Figure 4.7 summarizes the results of performing 25 trials of each sampling strategy for 9 different sampling rates ranging from 0.1 to 0.9. The ZP factor was set to 2 for each dimension. The presentation of recovery frequency is commonly used in the Compressed Sensing literature and is computed by establishing an arbitrary quality cutoff and counting the number of trials which exceeded the predefined threshold. A threshold of 40 dB was arbitrarily selected prior to running the experiment. Figure 4.7b depicts a plot of the average RSNR against the sampling rate for each sampling strategy. The error bars represent the sample standard deviation, or standard error, of the RSNR. These results indicate that the claims made by Hyberts *et al*. are valid and that the PG strategy out performs uniform random sampling.

4.5 Robustness to Noise

In order to investigate the robustness of a Compressed Sensing acquisition of the synthetic JPRESS data, the model data was scaled down to unity and complex white gaussian noise with standard deviation, σ, was added. The definition of SNR in the context of NMR is often debated and there exists no standardized way to measure it. Often the assessment is done in the frequency domain where some form of a ratio of the

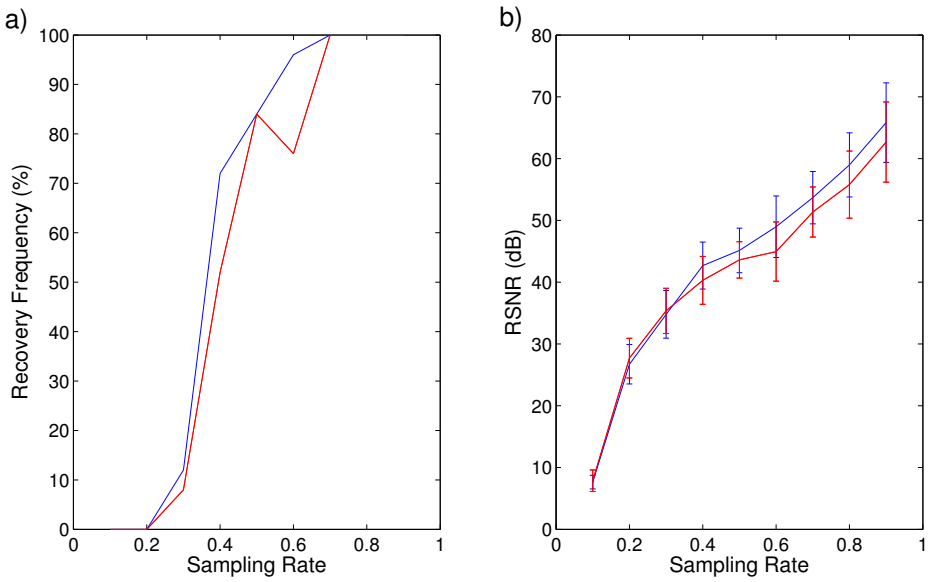

Figure 4.7: a) Recovery Frequency vs Sampling Rate b)Average RSNR (+/-SD) vs Sampling Rate

peak signal intensity is compared against a metric owing to a region of the spectrum devoid of NMR signals [35]. The issue was avoided by taking a qualitative approach, considering the noise as being low, moderate or high with standard deviations of 0.002, 0.01 and 0.05 respectively. The same methodology used in section 4.4 was applied except with the addition of noise into the measurements. The RSNR was calculated in the same way as previously. Only the PG sampling scheme was considered. Figure 4.8 plots the average RSNR against the sampling rate for the case of moderate noise.

The dashed red line in figure 4.8 is the RSNR calculated using the clean signal and the signal with low noise added. These results are in line with the theory presented in chapter 3. The RSNR metric suggests that a sampling rate of 0.2 produces better results than the fully sampled case. The issue here is that sparsity seeking algorithms

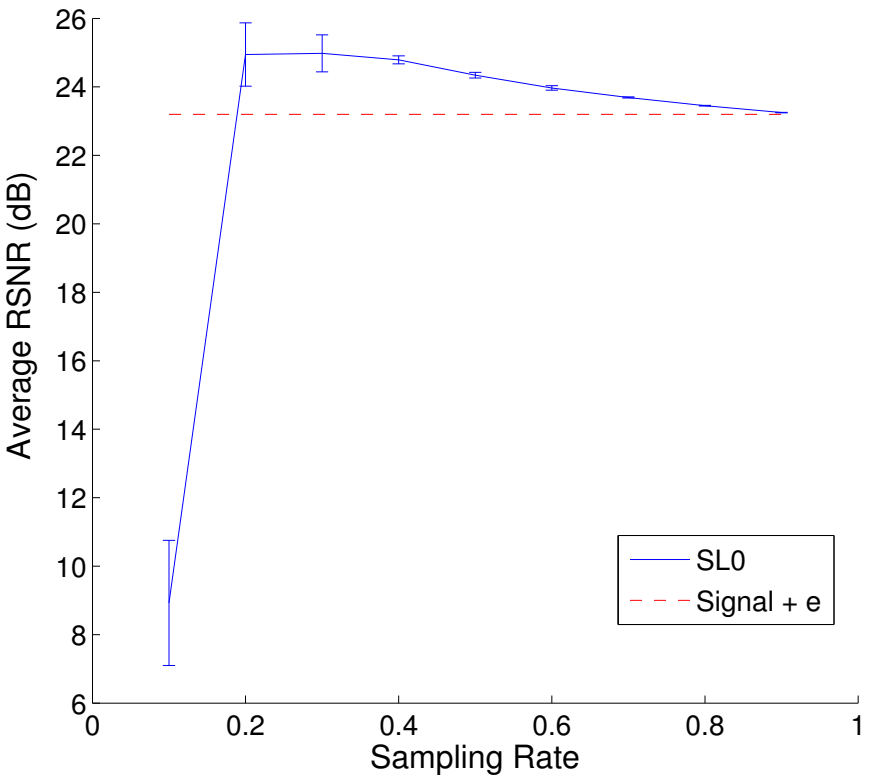

Figure 4.8: Average RSNR (+/-SD) vs Sampling Rate ($\sigma = 0.01$)

attempt to recover a sparse solution by comparing the measurement domain representation of the recovered signal against the original measurements. This is simply the spirit of basis pursuit in words. Basis pursuit de-noise and related approaches relax this fidelity constraint when the measurements themselves are contaminated with noise.

The denoise implementation of SL0, SL0-Denoise [36], was tested for the three noise conditions. In these experiments, the euclidean norm of the noise vector added to the measurements was passed as a parameter to the reconstruction algorithm which prevents projection in the case where the difference between the l_2 norm of

the measurements and solution vector is less than the l_2 norm of the noise vector. Although this is a contrived experiment, since in practice this parameter must be estimated by other means, it is worth investigating in the contexts of probing the limitations of applying Compressed Sensing to noisy systems. Figures 4.9b-d plot the results of applying the SL0-Denoise variant with the addition of low, moderate and high noise respectively. The dashed red line in all figures indicates again the baseline RSNR comparing the true signal with its noisy counterpart. For context, Figure 4.8a plots the reconstruction when no noise is added.

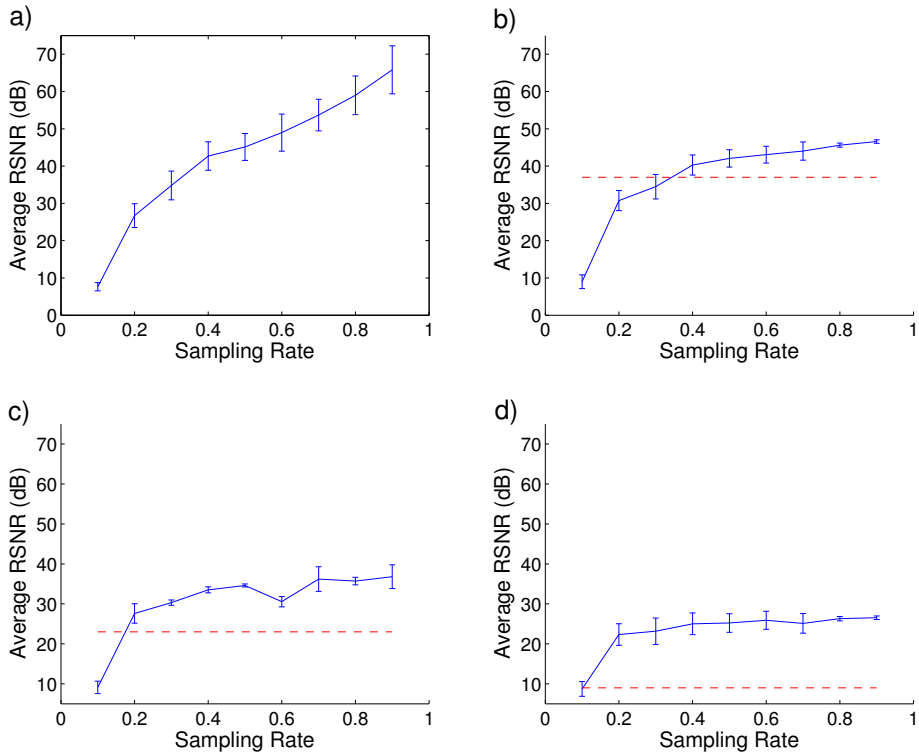

Figure 4.9: Average RSNR (+/- SD) vs Sampling Rate for a) No noise b) Low noise ($\sigma = 0.002$) c) Moderate noise ($\sigma = 0.01$) d) High noise ($\sigma = 0.05$)

4.6 Discussion

With regards to sampling strategies, the PG method out performed sampling from a uniformly random distribution. The sinusoidal modulation of the gap placement aims to place smaller gaps for the earlier time points and can not necessarily be justified in the language of Compressed Sensing. The justification is based off of the physical nature of NMR signals, such that they decay according to an exponential function. This entails that the majority of the signal content is front loaded near the beginning of acquisition, especially when noise unavoidably contaminates the measurements. This begs the question, why not uniformly acquire the first points only? This was addressed in [35] where the authors demonstrated that Compressed Sensing performs poorly in such a scenario. The reason is that NMR signals are not strictly sparse in frequency, but are more appropriately labelled as compressible in frequency. This owes to both the issue of spectral leakage as well as the Lorentzian lineshape. In the attempt to find a sparse solution, the tails of the lineshape are reduced, causing the growth of the amplitude of the signal in the time-domain. Linear prediction is a collection of methods that attempt to enhance resolution through the assumption that unmeasured points at the end of an acquisition could be conceived as being a linear combination of the points that precede. Compressed Sensing in its pure form fails as an implementation of linear prediction on the grounds stated. In [27], the authors exploited a sparsity enforcing algorithm to perform linear prediction with the additional constraint that the envelope of the signal should be consistent with the decay of the measured points.

The benefit of zero-padding in signal recovery was treated in detail in [33]. The authors argued that it is a naive attempt to reconcile the issue of spectral leakage in

the Compressed Sensing recovery of non-integral frequency complex exponential signals. Further investigation into the potential benefits and drawbacks of zero-padding are warranted, especially in the context of non-integral frequency decaying complex exponential signals.

By far the most troubling issue is the use of the RSNR metric in evaluating recovery success. One issue owes to the nature of JPRESS data that, due to imperfect water suppression and J coupling, there exists a vast dynamic range in the signals of interest. In the limiting case, where water is very poorly suppressed, the RSNR may be unrealistically high when in fact, the metabolite signals are poorly reconstructed.

The issue is further compounded in analyzing sparse recovery in the context of measurement noise. The experiments considered in this chapter had the benefit of full knowledge of both the true signal as well as the noise added. In practice, the retrospectively downsampled and subsequently recovered signal can only be compared to its fully sampled counterpart, which is unavoidably contaminated with noise. This leads to very pessimistic evaluations of the success of sparse recovery. This issue speaks to the problem of the vast gap between Compressed Sensing theory and its practical application.

The issue of RSNR was addressed in [37] where the authors concluded that the reconstruction of 23-Na images was adequate when only half the total measurements were considered. Like MRS, 23-Na imaging is notoriously insensitive and subject to considerable noise. Rather than evaluate the reconstruction based off a global metric like RSNR, the authors analyzed specific ROIs that contained the clinically relevant part of the signal. This assessment is a sort of parametric evaluation in which they concluded that the differences between the mean signal in the ROI in the fully and

undersampled cases were small enough to validate the success of Compressed Sensing in their application. This suggests that in attempting to evaluate the success of Compressed Sensing in a noisy regime, then perhaps the replicability of parameters, i.e. the things that motivate the measurements in the first place, should be used as a more appropriate benchmark.

Chapter 5

A Parametric Approach to Evaluating the Application of Compressed Sensing to In Vitro JPRESS Data

Chapter 4 demonstrated the feasibility of applying Compressed Sensing in the context of noiseless synthetic data. The results showed that when using PG sampling and introducing a moderate amount of zero-padding that the probability of accurately recovering the spectrum from its subsampled measurements is sufficiently high as long as approximately 50% of the samples are measured. The results were less clear when noise was incorporated into the measurements. A denoising variant of sparse recovery demonstrated improved results, but could only be implemented and verified with full knowledge of the true signal as well as the noise added. The focus of this

chapter is to employ a parametric approach to evaluate the success of sparse recovery of *in vitro* JPRESS data.

5.1 Materials and Methods

The JPRESS sequence was programmed in GE's proprietary platform dubbed Environment for Pulse Programming In C (EPIC) version 22. JPRESS was implemented by making two minor modifications to GE's existing PRESS sequence. The first modification was to introduce an incremental delay period between the first and second refocussing pulses and the second was to index the data in memory such that all averaged echoes for a given TE were stored separately. The sequence was installed on a 3T GE Signa750 MRI platform.

A 32-Channel GE head coil was used for all data acquisitions. The GE MRS phantom was placed in the centre of the coil and localized for positioning in the isocentre of the scanner bore. A series of scout images were acquired for spatial referencing. The sequence parameters are summarized in table 5.1. The voxel was placed in the centre of the MRS phantom and the automated prescan sequence was initiated. The stock prescan sequence calibrated the RF transmit and receiver gains, detects the centre frequency and optimizes field homogeneity. The field homogeneity was subsequently improved through manual adjustments. The flip angles of the water suppression pulses were also manually adjusted to maximize the efficacy of the global water suppression. The entire procedure took approximately 5 mins and the same values were used for all scans. Scan parameters and calibrations were downloaded to the system and the three scans were conducted consecutively. Finally, a non-water suppressed echo corresponding to the minimum echo time and with NEX = 8 was

acquired for post acquisition combination of the data. The raw data was transferred to a separate workstation where the individual coil data for each scan was combined with complex weights derived from the non-water suppressed data [38].

Parameter	Scan 1	Scan 2	Scan 3
TR	2 sec	2 sec	2 sec
TE(min)	31msec	31msec	31msec
SW_1	500Hz	500Hz	500Hz
SW_2	2000Hz	2000Hz	2000Hz
N_1	100	100	100
N_2	1024	1024	1024
Voxel Size	(3x3x3)cm^3	(3x3x3)cm^3	(3x3x3)cm^3
NEX	4	8	16
Scan Time	13:19 mins	26:36 mins	53:19 mins

Table 5.1: JPRESS Acquisition Parameters

The raw data for each scan was processed by ProFit. Usage of ProFit requires a license that is made available upon request through the Institute for Biomedical Engineering, ETH, Zurich. The current release is a Matlab only build and requires modification for handling raw data that did not originate from a Philips MRI scanner. The raw format is parsed by a Matlab script and returns a matrix containing the spectroscopic data as well as a data structure containing various details pertaining to the acquisition. Because the algorithm is based on prior knowledge, the experiment must match the simulated spectra with respect to acquisition parameters. The GE data was acquired with parameters (table 5.1) to match the simulated basis set. In order to analyze the GE data using ProFit, minor modifications were made to the manner in which the data matrix is passed through the referencing, phasing and quantification steps. No modifications were made to the main ProFit loop with the exception that various initializations were pulled out of the in the interest of speeding

up the hundreds of ProFit calls. The basis set was restricted from 20 metabolites to the 6 contained in the GE MRS phantom (table 4.1).

The process of retrospectively downsampling and reconstructing the data followed the same procedure outlined in chapter 4. For each sampling rate ranging from 0.1-0.9, 25 different sampling schedules were generated according to the PG method. Each reconstruction was referenced, phased and quantified. The denoising variant was not used.

5.2 Results

The ground truth was taken to be the metabolites ratio to Creatine, as it is a commonly used reference in MRS [14]. Just as the CRLB is commonly converted into a percentage of the estimated metabolite concentration, the metric used in this analysis was the absolute error of the recovered metabolite ratios,

$$Error_m = 100 \cdot \frac{\mid c_m/c_{Cr} - c_m^*/c_{Cr^*} \mid}{c_m/c_{Cr}} \tag{5.1}$$

where c_m/c_{Cr} is the ratio of the concentration of metabolite m to Creatine and c_m^*/c_{Cr^*} is the ratio of the concentration of the reconstructed metabolite m to the reconstructed Creatine concentration. Certainly, since it is a ratio, errors in the numerator and denomenator will serve to increase the overall error of the ratio; however, it is important to know how reliable such a measure is in reconstructed spectra.

Tables 5.2-7 display the mean ($\pm$SD) absolute error across the 25 trials for each sampling rate. Additionally, the maximum error is reported.

	Metabolite				
Sampling Rate	Naa	Ch	Glu	mI	Lac
0.1	2.34±2.0	7.33±3.8	15.0±9.7	15.5±10.1	13.4±4.6
0.2	0.71±0.6	0.73±0.4	4.38±2.0	2.34±1.7	3.73±2.6
0.3	0.97±0.5	0.49±0.3	1.49±1.6	1.87±1.4	3.61±2.3
0.4	0.60±0.4	0.36±0.3	0.77±0.7	1.40±0.9	1.73±1.5
0.5	0.53±0.3	0.38±0.3	0.72±0.5	1.41±1.0	1.16±1.1
0.6	0.33±0.3	0.34±0.3	0.57±0.4	0.80±0.5	1.07±0.6
0.7	0.25±0.2	0.26±0.2	0.46±0.5	0.72±0.6	0.82±0.5
0.8	0.25±0.2	0.25±0.2	0.49±0.4	0.49±0.35	0.70±0.5
0.9	0.14±0.1	0.14±0.1	0.29±0.3	0.33±0.3	0.28±0.2

Table 5.2: Mean (+/-SD) Metabolite Ratio Percent Errors, Scan 1 (NEX = 4)

	Metabolite				
Sampling Rate	Naa	Ch	Glu	mI	Lac
0.1	6.68	13.40	40.0	33.0	20.4
0.2	1.86	1.54	8.46	6.29	10.7
0.3	1.88	1.29	6.33	5.76	8.59
0.4	1.21	1.06	3.04	3.03	5.69
0.5	1.16	0.96	2.27	4.36	4.90
0.6	0.90	1.04	1.69	1.97	2.19
0.7	0.60	0.73	1.80	1.77	2.25
0.8	0.73	0.67	1.87	1.19	1.84
0.9	0.34	0.55	1.42	1.48	0.79

Table 5.3: Worst Case Metabolite Ratio Percent Errors, Scan 1 (NEX = 4)

5.3 Conclusions and Future Directions

The estimated metabolite ratios are in excellent agreement with the ground truth ratios corresponding to the fully sampled case. Due to the probabilistic nature of Compressed Sensing, reporting only the average case values can be misleading. When applying Compressed Sensing in a real world application, such as *in vivo* spectroscopy, there must be absolute certainty that the reconstruction is faithful to what might have been measured in the fully sampled case. If for instance, a physician agrees to

	Metabolite				
Sampling Rate	Naa	Ch	Glu	mI	Lac
0.1	4.76±2.3	7.82±3.7	13.4±9.4	14.0±9.7	11.8±6.4
0.2	0.60±0.5	0.85±0.7	3.00±1.5	2.04±1.4	4.45±2.5
0.3	0.63±0.5	0.50±0.3	1.49±1.3	1.15±0.9	2.98±2.4
0.4	0.42±0.3	0.36±0.3	0.63±0.5	1.10±0.9	1.64±1.1
0.5	0.52±0.3	0.33±0.2	0.60±0.5	0.75±0.5	1.24±0.8
0.6	0.52±0.7	0.29±0.3	0.50±0.5	1.08±1.6	1.43±1.2
0.7	0.33±0.2	0.23±0.2	0.40±0.3	0.56±0.6	0.88±0.5
0.8	0.27±0.2	0.20±0.2	0.28±0.3	0.43±0.4	0.47±0.4
0.9	0.16±0.2	0.15±0.1	0.23±0.2	0.22±0.3	0.35±0.3

Table 5.4: Mean (+/-SD) Metabolite Ratio Percent Errors, Scan 2 (NEX = 8)

	Metabolite				
Sampling Rate	Naa	Ch	Glu	mI	Lac
0.1	8.52	13.8	35.2	34.3	23.5
0.2	1.7	2.73	7.2	4.7	8.66
0.3	2.17	1.15	4.3	3.14	7.92
0.4	0.93	1.08	1.9	3.67	3.68
0.5	0.89	1.01	1.94	1.74	3.69
0.6	3.71	1.51	2.32	8.56	5.98
0.7	0.89	0.57	1.41	2.48	2.02
0.8	0.55	0.55	1.24	1.99	1.62
0.9	0.58	0.61	1.07	1.22	0.89

Table 5.5: Worst Case Metabolite Ratio Percent Errors, Scan 2 (NEX = 8)

perscribe MRS to assist in diagnosing the severity of a tumor, then a strong positive or negative result due to a poor reconstruction could have serious consequences for the patients treatment. In such a noisy regime, there is no way to infer on examination of a reconstructed spectrum that it is a "poor" reconstruction.

The results conveyed in this thesis make a compelling case for the application of Compressed Sensing to reduce the total acquisition time of 2D MRS; however, there is much work to be done. A major flaw with the methodolgy presented in this thesis is the lack of variety in the metabolites. The GE MRS phantom is itself a work horse,

	Metabolite				
Sampling Rate	Naa	Ch	Glu	mI	Lac
0.1	3.78±2.3	7.23±4.5	19.5±1.3	13.6±7.3	14.5±7.7
0.2	1.71±1.3	1.41±0.8	7.22±3.1	2.17±1.5	5.93±4.2
0.3	0.8±0.5	0.84±0.3	2.0±1.3	1.3±0.9	1.9±2.4
0.4	0.64±0.6	0.22±0.6	0.95±1.1	0.65±1.1	1.02±1.5
0.5	0.58±0.5	0.29±0.2	0.56±0.4	0.48±0.5	0.80±0.6
0.6	0.42±0.3	0.37±0.3	0.50±0.3	0.46±0.4	0.7±0.6
0.7	0.31±0.3	0.26±0.2	0.40±0.4	0.45±0.6	0.42±0.4
0.8	0.22±0.2	0.20±0.2	0.21±0.2	0.33±0.4	0.47±0.3
0.9	0.21±0.1	0.12±0.1	0.21±0.2	0.13±0.1	0.26±0.2

Table 5.6: Mean (+/-SD) Metabolite Ratio Percent Errors, Scan 3 (NEX = 16)

	Metabolite				
Sampling Rate	Naa	Ch	Glu	mI	Lac
0.1	7.75	18.0	44.1	27.3	27.8
0.2	4.94	3.61	15.1	5.95	4.9
0.3	2.44	2.64	4.55	4.75	5.58
0.4	1.81	0.65	2.51	1.53	2.48
0.5	1.53	0.72	1.62	1.83	2.55
0.6	0.97	1.03	1.15	1.48	2.28
0.7	1.16	0.65	1.44	2.49	1.55
0.8	0.70	0.58	0.95	1.93	1.09
0.9	0.49	0.37	0.79	0.47	0.74

Table 5.7: Worst Case Metabolite Ratio Percent Errors, Scan 3 (NEX = 16)

used week in and out in quality assurance and control, but it does a poor job of simulating the difficulties faced with *in vivo* MRS. These difficulties include patient motion, poor field homogeneity and contamination from nuissance signals such as lipids.

Additionally, only a single phantom was used. It is important to validate the technique in a wide variety of ratio levels. If a given disease process correlates with a low ratio for a certain metabolite, then this should be accurately captured in the reconstruction. Hyberts *et al*. examined this very scenario in the context of high resolution

NMR where they explored the probability of detecting weak peaks in Compressed Sensing acquisitions. The authors claimed that time-equivilent Compressed Sensing acquisitions, that is undersampling the data set but acquiring more averages such that the total duration is the same, can in fact improve the sensitivity of NMR.

Multidimensional spectroscopic imaging is also worth investigating. In 2012, Furuyama *et al*. demonstrated a fast spectroscopic imaging technique which accelerated the acquisition of J resolved images through Compressed Sensing [39]. Spectroscopic imaging combines the spatial encoding of conventional MRI with the spectral encoding of MRS to produce metabolite images. J resolved spectroscopic imaging data combines two spectral dimensions, as developed in this thesis, with at least two spatial dimensions, resulting in rich spatial-spectral information. The Compressed Sensing version was implemented using a fast spectroscopic imaging technique which couples the measurement of one spatial and one frequency dimension, leaving the remaing two to be stepped through incrementally. By randomly sampling the two indirect dimensions, the authors were able to exploit sparsity in the spatial-spectral domain. The concept may be taken a step further and beyond, spreading incoherence through all dimensions by exploiting even faster scanning techniques with added randomization, such as spiral imaging.

Bibliography

[1] Stephen F Keevil. Spatial localization in nuclear magnetic resonance spectroscopy. *Physics in medicine and biology*, 51(16):R579, 2006.

[2] John Pauly, Patrick Le Roux, Dwight Nishimura, and Albert Macovski. Parameter relations for the shinnar-le roux selective excitation pulse design algorithm. *Medical Imaging, IEEE Transactions on*, 10(1):53–65, 1991.

[3] Robin A De Graaf. *In vivo NMR spectroscopy: principles and techniques*. Wiley-Interscience, 2008.

[4] Helena Kovacs, Detlef Moskau, and Manfred Spraul. Cryogenically cooled probesa leap in nmr technology. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 46(2-3):131–156, 2005.

[5] DP Soares, M Law, et al. Magnetic resonance spectroscopy of the brain: review of metabolites and clinical applications. *Clinical radiology*, 64(1):12, 2009.

[6] R Bartha, DJ Drost, and PC Williamson. Factors affecting the quantification of short echo in-vivo1h mr spectra: prior knowledge, peak elimination, and filtering. *NMR in Biomedicine*, 12(4):205–216, 1999.

[7] M Albert Thomas, Noriaki Hattori, Masahiro Umeda, Tohru Sawada, and Shoji Naruse. Evaluation of two-dimensional l-cosy and jpress using a 3 t mri scanner: from phantoms to human brain in vivo. *NMR in Biomedicine*, 16(5):245–251, 2003.

[8] Varanavasi Govindaraju, Karl Young, and Andrew A Maudsley. Proton nmr chemical shifts and coupling constants for brain metabolites. *NMR in Biomedicine*, 13(3):129–153, 2000.

[9] Rolf F Schulte, Thomas Lange, Johannes Beck, Dieter Meier, and Peter Boesiger. Improved two-dimensional j-resolved spectroscopy. *NMR in Biomedicine*, 19(2):264–270, 2006.

[10] Jean-Baptiste Poullet, Diana M Sima, and Sabine Van Huffel. Mrs signal quantitation: a review of time-and frequency-domain methods. *Journal of Magnetic Resonance*, 195(2):134–144, 2008.

[11] Stephen W Provencher. Estimation of metabolite concentrations from localized in vivo proton nmr spectra. *Magnetic Resonance in Medicine*, 30(6):672–679, 1993.

[12] Ilya Kuprov, Nicola Wagner-Rundell, and PJ Hore. Polynomially scaling spin dynamics simulation algorithm based on adaptive state-space restriction. *Journal of Magnetic Resonance*, 189(2):241–250, 2007.

[13] SA Smith, TO Levante, Beat H Meier, and Richard R Ernst. Computer simulations in magnetic resonance. an object-oriented programming approach. *Journal of Magnetic Resonance, Series A*, 106(1):75–105, 1994.

[14] Rolf F Schulte and Peter Boesiger. Profit: two-dimensional prior-knowledge fitting of j-resolved spectra. *NMR in Biomedicine*, 19(2):255–263, 2006.

[15] JWC Van der Veen, R De Beer, PR Luyten, and D Van Ormondt. Accurate quantification of in vivo31p nmr signals using the variable projection method and prior knowledge. *Magnetic Resonance in Medicine*, 6(1):92–98, 1988.

[16] S Cavassila, S Deval, C Huegen, D Van Ormondt, and D Graveron-Demilly. Cramer–rao bounds: an evaluation tool for quantitation. *NMR in Biomedicine*, 14(4):278–283, 2001.

[17] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.

[18] Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.

[19] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on*, 52(12):5406–5425, 2006.

[20] Mark Rudelson and Roman Vershynin. Sparse reconstruction by convex relaxation: Fourier and gaussian measurements. In *Information Sciences and Systems, 2006 40th Annual Conference on*, pages 207–212. IEEE, 2006.

[21] Arian Maleki and David L Donoho. Optimally tuned iterative reconstruction algorithms for compressed sensing. *Selected Topics in Signal Processing, IEEE Journal of*, 4(2):330–341, 2010.

[22] David L Donoho. For most large underdetermined systems of linear equations the minimal ??1-norm solution is also the sparsest solution. *Communications on pure and applied mathematics*, 59(6):797–829, 2006.

[23] Rayan Saab, Rick Chartrand, and Ozgür Yilmaz. Stable sparse approximations via nonconvex optimization. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 3885–3888. IEEE, 2008.

[24] Ingrid Daubechies, Ronald DeVore, Massimo Fornasier, and C Sinan Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1):1–38, 2010.

[25] Hosein Mohimani, Massoud Babaie-Zadeh, and Christian Jutten. A fast approach for overcomplete sparse decomposition based on smoothed¡ formula formulatype=. *Signal Processing, IEEE Transactions on*, 57(1):289–301, 2009.

[26] Iddo Drori. Fast ? 1 minimization by iterative thresholding for multidimensional nmr spectroscopy. *EURASIP Journal on Advances in Signal Processing*, 2007, 2007.

[27] Alan S Stern, David L Donoho, and Jeffrey C Hoch. Nmr data processing using iterative thresholding and minimum¡ i¿ l¡/i¿¡ sub¿ 1¡/sub¿-norm reconstruction. *Journal of Magnetic Resonance*, 188(2):295–300, 2007.

[28] Krzysztof Kazimierczuk and Vladislav Yu Orekhov. Accelerated nmr spectroscopy by using compressed sensing. *Angewandte Chemie International Edition*, 50(24):5556–5559, 2011.

[29] Xiaobo Qu, Di Guo, Xue Cao, Shuhui Cai, and Zhong Chen. Reconstruction of self-sparse 2d nmr spectra from undersampled data in the indirect dimension. *Sensors*, 11(9):8888–8909, 2011.

[30] Daniel J Holland, Mark J Bostock, Lynn F Gladden, and Daniel Nietlispach. Fast multidimensional nmr spectroscopy using compressed sensing. *Angewandte Chemie*, 123(29):6678–6681, 2011.

[31] Sven G Hyberts, Alexander G Milbradt, Andreas B Wagner, Haribabu Arthanari, and Gerhard Wagner. Application of iterative soft thresholding for fast reconstruction of nmr data non-uniformly sampled with multidimensional poisson gap scheduling. *Journal of biomolecular NMR*, 52(4):315–327, 2012.

[32] Krzysztof Kazimierczuk and Vladislav Yu Orekhov. A comparison of convex and non-convex compressed sensing applied to multidimensional nmr. *Journal of Magnetic Resonance*, 2012.

[33] Michael Grant, Stephen Boyd, and Yinyu Ye. Cvx: Matlab software for disciplined convex programming, 2008.

[34] Hamed Firouzi, Masoud Farivar, Massoud Babaie-Zadeh, and Christian Jutten. Approximate sparse decomposition based on smoothed l0-norm. *arXiv preprint arXiv:0811.2868*, 2008.

[35] Marco F Duarte and Richard G Baraniuk. Spectral compressive sensing. *Applied and Computational Harmonic Analysis*, 2012.

[36] Sven G Hyberts, Koh Takeuchi, and Gerhard Wagner. Poisson-gap sampling and forward maximum entropy reconstruction for enhancing the resolution and sensitivity of protein nmr data. *Journal of the American Chemical Society*, 132(7):2145–2147, 2010.

[37] Sven G Hyberts, Scott A Robson, and Gerhard Wagner. Exploring signal-to-noise ratio and sensitivity in non-uniformly sampled multi-dimensional nmr spectra. *Journal of biomolecular NMR*, 55(2):167–178, 2013.

[38] Guillaume Madelin, Gregory Chang, Ricardo Otazo, Alexej Jerschow, and Ravinder R Regatte. Compressed sensing sodium mri of cartilage at 7t: preliminary study. *Journal of Magnetic Resonance*, 214:360–365, 2012.

[39] Mark Bydder, Gavin Hamilton, Takeshi Yokoo, and Claude B Sirlin. Optimal phased-array combination for spectroscopy. *Magnetic resonance imaging*, 26(6):847–850, 2008.

[40] Jon K Furuyama, Neil E Wilson, Brian L Burns, Rajakumar Nagarajan, Daniel J Margolis, and M Albert Thomas. Application of compressed sensing to multi-dimensional spectroscopic imaging in human prostate. *Magnetic Resonance in Medicine*, 67(6):1499–1505, 2012.