Math supplement: Magnetic resonance basics

Lars G. Hanson, February 2021

This mathematical appendix supplements the MRI introduction available in English and Danish via `https://drcmr.dk/Education#MRI_educational_material`.[1] Slides, software and YouTube videos are also provided there to improve the understanding.

1 The equation of motion for the nuclear magnetic moment

The hydrogen nuclei have a property called spin which make them behave like they are rotating, i.e. they have angular momentum (Danish: "bevægelsesmængdemoment" or "impulsmoment"). The fact that all nuclei appear to be spinning at *exactly* the same rate hints that there is more to it. Nuclear spin is a consequence of relativistic quantum mechanics, and is not quite rotation. For our purposes it is sufficient to note, however, that spin makes nuclei behave as particles rotating at the same rate.

Each hydrogen nucleus consists of a single positively charged proton. The current loop associated with a rotating charge distribution is expected to cause magnetism, and that is indeed what happens for the hydrogen nuclei. They each have a magnetic moment $\boldsymbol{\mu}$, which is proportional to their angular momentum $\mathbf{J}$ that is along the spin axis.

$$\boldsymbol{\mu} = \gamma \mathbf{J} \tag{1}$$

The constant of proportionality is the "gyromagnetic ratio" which differs between nuclear species and is $\gamma = 42$ MHz/T for hydrogen nuclei. We will immediately convert Hz to an angular frequency by multiplying γ with 2π: $\gamma = 263 \cdot 10^6$ s^{-1}/T. Depending on context, γ may be defined to include the factor of 2π in which case the unit is s^{-1}/T. Similarly, the Larmor frequency may be given in Hz or in radians per second (s^{-1} since radians are implicit). This may be confusing at first, but it is typically not too difficult to resolve. You just need to make sure that γ and the Larmor frequency appear with the appropriate units whenever they matter.

We will now consider how the magnetic dipole moves in a magnetic field. From classical mechanics it is known that a change of the angular momentum is generally related to

[1] For DTU courses 22481 and 22485, the mathematical description of magnetic resonance (section 4) is not essential, whereas the math in the first sections are (relaxation in 22481, precession and relaxation in 22485). The understanding of magnetic resonance is essential in both courses. Section 5 concerning spectroscopy and imaging is a topic of course 22506. See `https://www.cmr.healthtech.dtu.dk/education/mr-courses-and-their-connection`

the torque $\boldsymbol{\tau}$ ("kraftmoment" in Danish): $d\mathbf{J}/dt = \boldsymbol{\tau}$. In this case, the torque is the "twist" caused by the magnetic field $\mathbf{B}$ on the magnetic dipole:

$$\frac{d\mathbf{J}}{dt} = \boldsymbol{\mu} \times \mathbf{B} \tag{2}$$

When merging the equations above, we get a general formula describing how the magnetic moment of a rotating particle will change when subject to a magnetic field:

$$\frac{d\boldsymbol{\mu}}{dt} = \gamma \boldsymbol{\mu} \times \mathbf{B} \tag{3}$$

Since the change of $\boldsymbol{\mu}$ is orthogonal to itself and to $\mathbf{B}$, this equation describes rotation of $\boldsymbol{\mu}$ around $\mathbf{B}$. This is true for the nuclei individually (or strictly speaking for their $\boldsymbol{\mu}$ expectation values, if you are into quantum mechanics). Provided the nuclei feel the same field, the equation also applies to the total magnetization of all nuclei, $\mathbf{M} = \sum_i \boldsymbol{\mu}_i$ (also called the net magnetization or "nettomagnetiseringen" in Danish). Check that this latter claim is true!

2 Precession

Before exploring the magnetic resonance phenomenon, we consider how the magnetization of a nucleus will behave in a static magnetic field $\mathbf{B}_0 = (0, 0, B_0)$, where the z-axis has been defined to be along the static field $\mathbf{B}_0$. Using the equation above for this special case, we have

$$\frac{d\mu_x}{dt} = \gamma \mu_y B_0, \qquad \frac{d\mu_y}{dt} = -\gamma \mu_x B_0, \qquad \frac{d\mu_z}{dt} = 0 \tag{4}$$

These coupled differential equations, can be solved in several ways, and we choose a particularly elegant one: Looking at equation (3), we may realize that the magnetization will rotate around the magnetic field. Since rotation can be described by complex exponentials, we choose to represent the transversal magnetization as a complex number, even though it actually is a 2D vector. It will turn out to be convenient to combine the x and y components of $\boldsymbol{\mu}$ into a complex transversal magnetization $\mu_{xy} \equiv \mu_x + i\mu_y$. Using equation (4) and this definition, we do a bit of rewriting (verify it!):

$$\frac{d\mu_{xy}}{dt} = -i\gamma B_0 \mu_{xy}, \qquad \frac{d\mu_z}{dt} = 0 \tag{5}$$

The shift to a complex notation reduced the three coupled equations to two uncoupled first-order differential equations. We recognize the Larmor frequency $\omega_0 = \gamma B_0$ in these equations which are easily solved:

$$\mu_{xy}(t) = \mu_{xy}^{t=0} \exp(-i\omega_0 t), \quad \mu_z(t) = \mu_z^{t=0} \tag{6}$$

We are free to choose time zero and the initial conditions as we wish.

Since multiplication by a factor $\exp(i\phi)$ rotates a complex number by an angle ϕ in the complex plane, we see that this solution to the equations of motion (5) indeed describes clockwise rotation of the transversal magnetization with a frequency ω_0.

3 Relaxation

We have found the time evolution of the nuclear magnetization in a constant magnetic field, but we still need to include relaxation in the description. This comes about since the individual nuclei interact magnetically so that they feel field fluctuations around the mean. This makes the total transversal magnetization M_{xy} decay exponentially towards zero with a time constant T_2. Similarly, the total longitudinal magnetization M_z is known to relax exponentially back towards equilibrium $\mathbf{M}_0 = (0, 0, M_0)$ on a time scale T_1. We may describe the evolution as follows:

$$\begin{aligned} M_{xy}(t) &= M_{xy}^{t=0} \exp(-i\omega_0 t) \exp(-t/T_2) \\ M_z(t) &= M_z^{t=0} \exp(-t/T_1) + M_0(1 - \exp(-t/T_1)) \end{aligned} \tag{7}$$

The first equation describes how the precessing transversal magnetization decays away. For the intensity of an image acquired at time TE after excitation (the "echo time"), the phase factor is insignificant, $I(\mathrm{TE}) \propto |M_{xy}(\mathrm{TE})| = |M_{xy}^{t=0}| \exp(-\mathrm{TE}/T_2)$. The first term of the second equation describes how the "memory" of the initial longitudinal magnetization $M_z^{t=0}$ decays away on a timescale T_1 while the second term describes how "fresh" longitudinal magnetization is created on the same time scale so that the magnetization approaches equilibrium after excitation. Please prove that this formulation is equivalent to the following: The deviation of the longitudinal magnetization from equilibrium, decays away exponentially with a time constant T_1.

The evolution of the magnetization could alternatively be found by directly solving modified differential equations similar to equation (5) with added relaxation terms:

$$\frac{dM_{xy}}{dt} = -i\omega_0 M_{xy} - M_{xy}/T_2, \qquad \frac{dM_z}{dt} = -(M_z - M_0)/T_1 \qquad \text{(Check!)} \tag{8}$$

4 Magnetic resonance

We now explore the magnetic resonance phenomenon which provides a way to rotate the magnetization away from equilibrium using weak oscillating magnetic fields. We ignore relaxation since the involved bursts of radio waves are so short that relaxation during excitation is insignificant. Hence the description below is equally valid for the magnetic moment $\boldsymbol{\mu}$ of a single nucleus, and for the total magnetization $\mathbf{M}$ of nuclei that all feel the same external magnetic fields (such a group of nuclei is called "an isochromate").

MR experiments involve a strong polarizing field, $\mathbf{B}_0$, and a much weaker orthogonal near-resonant RF field, $\mathbf{B}_1(t)$, oscillating in the xy-plane (we can ignore the electric contribution to the RF field in this connection). Again, we combine the x and y components and define a complex magnetization $\mu_{xy} = \mu_x + i\mu_y$ and a similar complex field $B_{xy} = B_x + iB_y$. Rewriting equation (3), we get

$$\frac{d\mu_{xy}}{dt} = -i\gamma(\mu_{xy}B_z - \mu_z B_{xy}), \qquad \frac{d\mu_z}{dt} = \gamma \,\mathrm{Im}(\mu_{xy}^* B_{xy}) = \frac{-i\gamma}{2}(\mu_{xy}^* B_{xy} - \mu_{xy} B_{xy}^*) \tag{9}$$

The asterisk (*) denotes complex conjugation. The motion of $\boldsymbol{\mu}$ is dominated by a rapid clockwise precession around $\hat{\mathbf{z}}$ for positive γ since $|B_z| = B_0 \gg |B_{xy}|$. The case of a

constant amplitude RF field is of special interest and is discussed in detail. We define the x-axis to be along the oscillating field generated by a loop coil: $\mathbf{B}_1 \equiv 2B_1 \cos(\omega t)\hat{\mathbf{x}}$. Using the definition of B_{xy} and Euler's formula, this is equivalent to a complex field

$$B_{xy} = B_x + iB_y = B_1(\exp(i\omega t) + \exp(-i\omega t)) \tag{10}$$

The component $B_1 \exp(-i\omega t)$ rotating the same way as the precessing magnetization will influence it significantly, while the counter-rotating field $B_1 \exp(i\omega t)$ can safely be ignored since it is roughly $2\omega_0$ away from resonance (it causes a slight wiggling only). This approximation is valid for fields oscillating near the Larmor frequency, and we only care about such. Hence, half of the field generated by a loop coil is wasted (the counter-rotating component), and we generally prefer circularly polarized coils generating and receiving fields that rotate together with the magnetization.

The equations of motion are expressed in the rotating frame of reference following the B_{xy}-field rotation by defining the slowly varying quantities $\tilde{\mu}_{xy} = \mu_{xy} \exp(i\omega t)$ and $\tilde{B}_{xy} = B_{xy} \exp(i\omega t) \simeq B_1$. The approximation in the last equation was justified above. Expressed in terms of those quantities, equation (9) becomes

$$\frac{d\tilde{\mu}_{xy}}{dt} = i(\omega - \omega_0)\tilde{\mu}_{xy} + i\gamma B_1 \mu_z \tag{11}$$

$$\frac{d\mu_z}{dt} = \gamma \mathrm{Im}(\tilde{\mu}^*_{xy}\tilde{B}_{xy}) = -\gamma \tilde{\mu}_y B_1 \tag{12}$$

The y-component of the transversal magnetization in the rotating frame of reference $\tilde{\mu}_{xy}$ is here denoted $\tilde{\mu}_y \equiv \mathrm{Im}(\tilde{\mu}_{xy})$. For the special case of a resonant field, $\omega = \omega_0$, the equations reduce to

$$\frac{d\tilde{\mu}_{xy}}{dt} = i\gamma B_1 \mu_z \tag{13}$$

$$\frac{d\mu_z}{dt} = -\gamma \tilde{\mu}_y B_1 \tag{14}$$

If the first equation is split in real and imaginary parts, we see that we are back to a set of equations very similar to (4). The solution is therefore also similar, now being precession of $\tilde{\boldsymbol{\mu}}$ around $\tilde{\mathbf{B}}_1$ (rather than of $\boldsymbol{\mu}$ around $\mathbf{B}_0$). Consequently, when a constant RF field is present, the dynamics in the rotating frame is found to resemble those in the stationary frame in the absence of RF field. The only difference is that the precession is now around an axis in the transversal plane.

To summarize: In absence of radio waves, the magnetization precesses around the static magnetic field $\mathbf{B}_0$ at the frequency γB_0. Changing to the rotating frame of reference, this precession "disappears", and the magnetization becomes stationary. The dynamics become complex in the stationary frame, when a resonant rotating radio wave field $\mathbf{B}_1$ is added. But in the rotating frame, this field is stationary, and the magnetization will simply precess around it at a low frequency γB_1. This is nuclear MR, which differs from compass MR only since the nuclei precess rather than vibrate. The difference is due to the nuclei not only being magnetic, but also having spin (angular momentum). The shift to a complex notation, and to the rotating frame of reference, is so convenient that even the sampled raw signal is saved by the scanner in a corresponding format.

Exercise: Rewrite equations (8) in a frame of reference rotating at frequency ω_0. Solve!

5 From Spectroscopy to Imaging to Spectroscopic Imaging: Phase rolls and the Fourier transform

It is tempting (and often useful) to explain imaging in terms of frequency encoding, which is easy to understand, but has limited scope, since it only explains sequences with constant readout-gradients sufficiently, i.e., 1D or radial sequences. The k-space approach is much more general, and will be introduced here, and even extented to (k,t)-space which is central for spectroscopic imaging. Math and explanations are mixed, so keep reading, even if parts are challenging. The reader is assumed to have a basic understanding of MR spectroscopy and classical MR concepts[2], e.g. as described in introductory sections of notes at `http://eprints.drcmr.dk/37/`.

There are many ways of introducing imaging and slightly different approaches are chosen here and in the mentioned MR introduction, which may cause slight confusion (or clarity). The mentioned notes only describe gradients as creating phase rolls during imaging, whereas gradients are here also described as refocusing pre-existing conceptual phase rolls (details below). The latter description is more accurate and general, but also more challenging to understand. To keep the explanation simple, the course notes therefore implicitly assume k-space symmetry, which is not always valid. This text is a supplement.

The starting point for the coming discussion is basic MR and spectroscopy, which is assumed familiar (see course notes). The first parts explain how the Fourier transform works, and is general for any kind of spectroscopy and imaging. The word "scanner" will in this context be used indiscriminately for NMR spectrometers and imaging systems. Only proton MR is considered for convenience, but the principles are general.

1. Consider excitation by a 90° pulse converting longitudinal magnetization into transversal magnetization at time $t = 0$. We can now record an oscillating signal curve (FID) reflecting precession. It is first discussed how the Fourier transform converts the FID into a spectrum, and this insight is used to understand imaging.

   We define the vector $\mathbf{M} = (M_x, M_y, M_z)$ to be the magnetization of the sample. As always, the z-axis is chosen along the B_0-field. The transversal magnetization (M_x, M_y) is a 2D vector, and such vectors can alternatively be represented as complex numbers. We prefer to use the latter for MR to simplify math, since rotation is compactly described using complex numbers. We define the complex transversal magnetization $M_{xy} = M_x + iM_y$. Clockwise rotation of a complex number by an angle ϕ around origo is simply described by multiplication with a factor $\exp(-i\phi)$ (never mind why – you may choose to think of it simply as a notation for rotation). Clockwise precession at some frequency ω_m can therefore be conveniently expressed as a complex phase factor $\exp(-i\omega_m t)$ that is multiplied onto the initial transversal magnetization $M_{xy}(t = 0)$ created by excitation. ω_m is

[2] Classical MR descriptions are consistent with correct QM descriptions, as described in *Is Quantum Mechanics necessary for understanding Magnetic Resonance?* Concepts in Magnetic Resonance Part A, 32A(5), 2008. There is additional detail in a new book chapter partially available at `http://drcmr.dk/MR`

here the metabolite precession frequency. The figure below shows the transversal magnetization at different points in time following excitation. The rotation that is seen in the transveral (or complex) plane reflects precession. The observation period is here chosen so short that relaxation plays no role, but on a longer time scale, the arrow length will decrease on a time scale $T_2{}^*$.

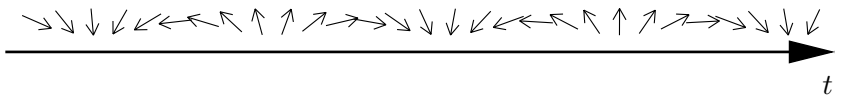

"Phase" in the context of MR typically refers to the direction of the transversal magnetization, i.e. the angle that changes due to precession in the figure above. Precession causes a linear phase variation $\phi = \omega_m t$ along the time axis, i.e., a regular rotation called a **phase roll** as illustrated above.

2. The signal recorded after excitation is proportional to the transversal magnetization and will reflect precession and signal decay.[3] It is frequency-shifted in the scanner's receiver by the demodulation frequency that is typically set to the Larmor frequency of water. The demodulation simply slows all recorded precession down by the Larmor frequency of water. This results in the transversal magnetization appearing to be recorded in a frame of reference rotating at the water frequency (see mentioned course notes). The metabolite signals will each oscillate as illustrated above, but at different frequencies determined by the molecular structure (the chemical shift). This rotation is reflected in the recorded samples that are stored as complex numbers changing as shown above (the FID). Due to demodulation at the water frequency, only the water signal will not oscillate:

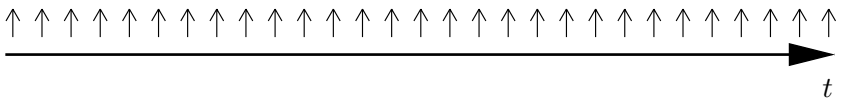

Consider the integral $\int S(t)\,dt$ of the signal over time, i.e. the sum of all the complex numbers (representing magnetization vectors) illustrated above. It will be proportional to the non-oscillating water content, since the other metabolite signals oscillate and therefore contribute insignificantly to the integral (all phases are almost equally present for those). The total signal is a sum of metabolite contributions, $S(t) = \sum_m S_m(t=0)\exp(-i\omega_m t)$, where the term contributed by water has $\omega_m = 0$ due to demodulation. To separate out another signal component than water, we need to compensate its precession before integrating. Back-rotation to detect a frequency ω is done by multiplication of the signal $S(t)$ by $\exp(i\omega t)$ before integration, i.e. $\int S(t)\exp(i\omega t)\,dt$. Remember that the signal from a particular metabolite, $S_m(t)$, is itself proportional to $\exp(-i\omega_m t)$ where ω_m is the metabolite frequency. When ω equals ω_m, the two exponentials cancel, and the integral is proportional to $S_m(t=0)$ and therefore to the metabolite content.

[3] the constant of proportionality is complex since the coil is sensitive to the change of the magnetization rather than the magnetization itself, but that is a detail in this context.

In summary, the precession causes a linear phase variation $\phi = \omega_m t$ along the time axis, i.e. a **phase roll** (the wavy pattern above). This can for each metabolite be refocused by multiplication with a complex exponential in the Fourier integral, which simultaneously dephase other metabolite signals, if present. Integration quantifies the refocused component, and eliminates other contributions. This insight into the inner workings of the Fourier transform is very useful in many contexts, including imaging. The concept of phase rolls is equally important.

The signal $S(t) = \sum_m S_m(t = 0) \exp(-i\omega_m t)$ was here expressed as a sum over metabolites oscillating at discrete frequencies. More generally, a continuum of frequency components may be present, e.g., if there is field variation. The total signal is then alternatively described as an integral over a frequency distribution, $S(t) = \int S(\omega, t = 0) \exp(-i\omega t)\, d\omega$.

In brief, chemical shift gives rise to a phase roll along the time axis. This is refocused (un-rolled) by the complex exponential in the Fourier integral when the spectrum is calculated. The integral evaluated at a metabolite frequency (the peak height) will be proportional to the metabolite content, but will also depend on the relaxation time (remember that only the beginning of the FID was shown above).

3. Next, imaging is considered, i.e., measuring the spatial distribution of transversal magnetization, $M_{xy}(\mathbf{r})$. For this we need to introduce the concept of gradients. Scanners are equipped with three "gradient coils" that cause linear field variation along the x, y and z direction, respectively. A combination of gradient fields is itself just a linear field variation in some slanted direction. Using the gradient coils, we can introduce linear field variations in any direction we wish. The gradient field adds onto whatever field variation that may already be present. The latter variation is assumed zero in the following (perfect shim), and we will also assume initially that only water gives significant signal.

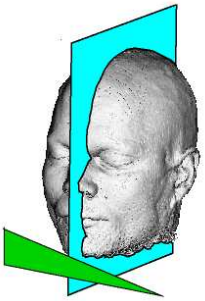

A gradient from the patient's left to right side. All nuclei within the shown sagital slice experience the same field, and therefore have the same Larmor frequency.

4. During excitation and before any gradient is applied, the nuclei will be rotated equally by the RF field (assumed homogeneous), and the local transversal magnetization therefore ends up in the same direction everywhere. We say that it is "in phase" across the sample, and the signal immediately after excitation will therefore be strong and reflecting the total proton content, since all contributions to the total magnetization are aligned.[4] The corresponding recorded signal contains no spatial information, however, and we therefore need gradients to provide spatial discrimination. Slice selection is a simple and well-known way to reduce the 3D imaging problem to a 2D problem: We turn on the gradient shown above during excitation. Only nuclei on resonance are affected significantly by radio waves, so by sending such at the Larmor frequency of the nuclei in the shown sagital slice, it is possible to excite only these. We choose to rotate the nuclei within the slice 90 degrees, as before, and switch the slice selection gradients off at a time after excitation where all nuclei in the slice are "in phase", i.e., point in the same transversal direction. With the B_0-field along the body, they may at time zero point toward the nose, for example, and they will subsequently precess in phase until we apply more gradients. Much more can be said about slice selection, but the focus is here on other aspects of imaging, so we hurry on.

5. Right after excitation of the slice, the local contributions to the magnetization are aligned as described above. When another gradient is applied in-plane after excitation, there will be a linear field variation across the patient, so the nuclei will start precessing at different frequencies along the direction of the gradient. We can apply this gradient in the direction from neck to nose, for example, and after excitation of the sagital slice. The nuclei will now dephase in a controlled way: Neighboring nuclei feel almost the same field whereas nuclei far apart are rotated quite differently, and end up in opposite directions, for example. To be specific, a spatial phase roll is accumulated, i.e., a linear phase variation along the direction of the applied gradient. This is similar to the temporal phase roll $\exp(-i\omega t)$ described above for spectroscopy.

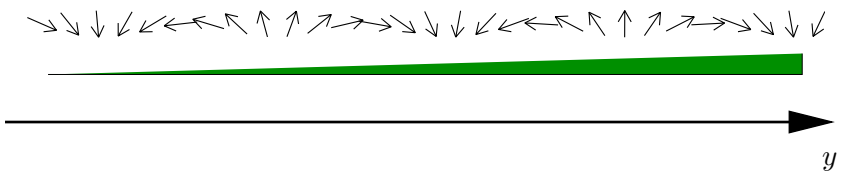

The spatial phase variation can be described as $\exp(-i2\pi k_y y)$ for a gradient applied along the y-axis, e.g., a linear field variation from neck to nose. The extra constant 2π only appears to enforce a simple relation between the wavelength λ_y of the phase roll, and the constant $k_y = 1/\lambda_y$. We notice that the phase factor is indeed periodic in y and unchanged whenever λ_y is added to y (2π rotation corresponds to 360°).

[4] Not all nuclei are aligned but the magnetizations from different local nuclear ensembles having essentially experienced the same field are aligned, i.e. the local isochromate net magnetizations.

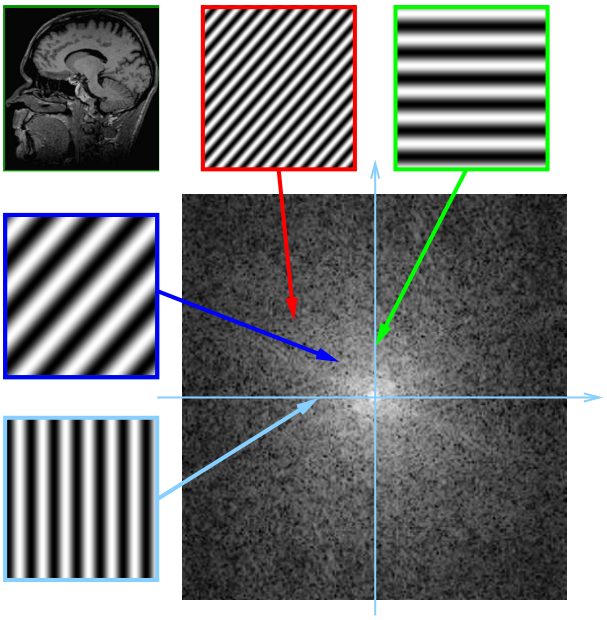

Figure 1: The structure of k-space which is a (k_x, k_y)-coordinate space where each point correspond to a particular phase roll pattern, here indicated by color coded phase images (phase angle mapped to color). The background color in the large image indicates the measured RF signals (Fourier coefficients $W(\mathbf{k})$) for a particular patient (corner) after the phase roll patterns are refocused one by one using gradients. k-space is a book-keeping tool used to keep track of the phase roll patterns needed to do imaging: We need to measure the MR signal $W(\mathbf{k})$ for each of the phase roll patterns $\mathbf{k}$ corresponding to the shown central region of k-space. The more of k-space we cover, the more detailed will the resulting MR image be.

A phase roll is characterized uniquely by its wave length and the direction of it.[5] A general phase roll in any direction can be described as $\exp(-i2\pi\mathbf{k}\cdot\mathbf{r})$, i.e,. a phase (precession angle) that varies linearly in direction of a vector $\mathbf{k}$, and with a wavelength $1/|\mathbf{k}|$. The vector $\mathbf{k}$ thus describes the phase roll pattern created by application of gradients uniquely. Right after excitation, $\mathbf{k}$ is zero since no nuclei are yet dephased by gradients. When gradients are applied, $\mathbf{k}$ changes since the phase roll changes.

As long as the gradient field is active, the phase roll will change, and therefore also the $\mathbf{k}$-vector. Specifically, the wavelength will become progressively shorter if a constant gradient is applied after excitation (the wavelength of the phase roll is the distance between repetitions of the wavy pattern). To see the phase roll forming, go see `https://www.youtube.com/watch?v=qXhQhgvpRU0`.

[5] Except for a spatial shift controlled by a complex coefficient later multiplied onto the phase roll.

For a homogeneous substance, application of a gradient that causes several full phase rolls (rotations) along the gradient axis will leave almost equally many nuclei pointing in all transversal directions. The coil measures the contribution from all nuclei, i.e. the spatial integral, which plays the same role as the temporal integral experienced earlier. Hence the signal is suppressed (spoiled) by phase rolls, and it may seem that little is gained in terms of getting spatial information. The signal, however, depends on the structure of the object, and will remain strong if the structure matches the phase roll (striped patient). Specifically, what is measured after a phase roll is created, is the corresponding spatial Fourier component, i.e. the "stripedness" of the patient in the scanner for a particular stripe pattern (phase roll). The stripe pattern is selected via the gradient direction and the duration of the gradient, if it is a constant gradient. More generally, the stripe pattern depends on the entire gradient (and RF) history since excitation.

To understand what is meant by "stripedness", we note that phase roll patterns appear striped when the phase of the transversal magnetization is color-coded on a gray-scale so that, e.g., x-magnetization is colored black, and $-x$-magnetization is colored white. The mentioned "stripe patterns" above should be interpreted in that sense, i.e., as phase roll patterns, which can indeed graphically be represented as wavy stripe patterns.

6. On the way to establishing that the stripedness is indeed measured, we now turn to the Fourier theorem, which roughly can be formulated as follows: Any spatial function can be expressed as a weighted sum of phase roll patterns, or in our case, any complex image can be represented as a weighted sum of stripe patterns (magic, but a simple argument is given towards the end of the text). This is specifically true for the transversal magnetization distribution resulting from excitation (before application of any gradients), $M_{xy}(\mathbf{r}) = \sum_{\mathbf{k}} W(\mathbf{k}) \exp(i2\pi\mathbf{k}\cdot\mathbf{r})$, where $W(\mathbf{k})$ are the weighting factors for phase rolls identified by $\mathbf{k}$. The excitation may possibly be followed by some contrast preparation, e.g. a relaxation period that will change M_{xy} and therefore the weights, but the general expression above still applies. The equation expresses that whatever the spatial distribution of magnetization is, it can be perceived as a sum of phase rolls. These are pre-existing in the sense that they are not created by imaging gradients. They arose as conceptual, but can be perceived as very real, which the following discussion will show.

7. In all relevant cases, the sum above is well approximated by a finite sum of low-frequency phase roll patters, i.e., a sum over $\mathbf{k}$-vectors in the central region of k-space. The Fourier theorem tells us that measuring $W(\mathbf{k})$, i.e., the stripedness of the patient for all relevant stripe patterns (phase rolls) is as good as measuring $M_{xy}(\mathbf{r})$ itself. If we know $W(\mathbf{k})$ we can get to the latter by calculating a weighted sum of stripe patterns, where the measured stripedness is chosen as the weighting factor. In other words, the similarity of the patient to a particular stripe pattern is chosen as the weighting factor of that stripe pattern when the image is calculated.

We now just need to establish that we can indeed measure the stripedness $W(\mathbf{k})$ of the patient, and we are almost there: Initially it was argued that a phase roll along the time axis (caused by chemical shift) can be refocused by means of multiplication with an opposite phase roll, and that subsequent integration over time gives a measure of the corresponding metabolite content. Similarly, the Fourier expansion above describes the magnetization distribution that we want to image, as a weighted sum of phase rolls that we can each refocus one at a time to measure the coefficients $W(\mathbf{k})$. We have learned that gradients create phase rolls, but similarly they can refocus the conceptual pre-existing phase rolls appearing in the Fourier expansion.

The challenge of separating metabolite signals in spectroscopy is replaced by a challenge of separating out spatial components mixed since the receiver coil measures only the total transversal magnetization within the coil volume. When analyzing the measurements, we single out spectral and spatial components using the same trick: In spectroscopy, the phase rolls that were created by chemical shift were removed mathematically. In imaging, the conceptual phase rolls were removed with gradients to measure specific Fourier components, but need to be reintroduced to form images: Once the Fourier coefficients $W(\mathbf{k})$ are measured by refocusing the conceptual phase rolls in the Fourier expansion individually using gradients, we use that same Fourier expansion to calculate the image. This is called image reconstruction, and is seen to involve multiplication with complex exponentials and integration, as for spectroscopy.

8. The principles above were described for slice-selective 2D imaging to ease visualization. In the math, however, we did not use this, and the Fourier theorem applies in more dimensions. The expansion to 3D imaging is therefore mathematically straight-forward, and the phase rolls are then 3D. Instead of stripe patters, the corresponding graphical representation will be stacks of equidistant planes separated by the wavelength. The planes can be oriented in any direction, and each such stack are represented by a single point in 3D k-space (k_x, k_y, k_z).

   We can conclude even more, in fact. Spectrospic imaging is a measurement providing a full spectrum for each position in an imaged volume (e.g. a slice). In principle, this requires a kind of movie, where the frame rate is so high that the precession can be followed, and the metabolites therefore be identified. An FID per image voxel is in other words required. Acquiring such images at slightly different sampling times indeed allows for spectroscopic imaging, but an extension of the arguments above shows that we have much more flexibility. We have seen that chemical shift has the same effect as a gradient creating phase rolls in space: It creates a phase roll in time, and thus acts as a temporal gradient. We have also learned that a sum of gradients is itself just a gradient, and this is true, even if it is a gradient in a mix of spatial and temporal dimensions. With no further argument, we can conclude that if we manage to map the signal not only in k-space, but also the time evolution, i.e. $W(k_x, k_y, k_z, t)$, then we can reconstruct

spatially resolved spectra $S(x, y, z, \omega)$ by adding together phase rolls in a 4D space. We need not cover this $(\mathbf{k}, t)$-space in a way that invites calculation of a movie, as described above. We just need to sample it critically, i.e. fulfill the Nyquist criterium everywhere. We kan limit that space to fewer dimensions by doing slice selection eliminating one spatial dimension, e.g., x for a sagital image. The remaining dimensions, (k_y, k_z, t), can for example be covered slicewise, so one slice of $(\mathbf{k}, t)$-space, (k_z, t), is covered per excitation. This is known as echo planar spectroscopic imaging.

9. The Fourier theorem was crucial in the argument above, and it may seem unlikely that any image can indeed be expressed as a weighted sum of stripe patterns corresponding to oscillatory functions, each oscillating around 0. You can possibly convince yourself by following this train of thought: Imagine gradually adding many such stripe patterns together, where each of them is chosen to be bright in one particular point, e.g., the middle, but where wavelength and orientation of stripes are chosen randomly. The more stripe patterns you add together, the brighter will the chosen point be, whereas the intensity in all other points will fluctuate randomly, depending on which stripe pattern you happen to add last. You will therefore end up with a bright spot, and a backgrond that is dark *in comparison*. Since you have now constructed an image of any single point by calculating a weighted sum of phase roll patterns, you can do much more. You can construct any image point by point, and add all the mixes for different points together subsequently. You then have a full image constructed as a sum of phase roll patterns.