Chapter 4

Helioseismic Inverse Problems

In this section some of the general properties of inversion methods will be described. Most of the results presented here are from published papers or papers included in the appendix of which I have been co-author. In addition to the descriptions of the methods presented in those papers and those written by others (see the references in the included papers), a 'manual' for the 1- and 1.5-dimensional inversion package which I have written (Schou 1993c), contains some details on how various things can be implemented.

While inversions can be performed to infer other aspects of the solar interior than the differential rotation, I have mainly been concerned with this problem. The rotation problem is particularly simple, as it is to a very good approximation a linear problem, unlike for example, determining the sound speed in the solar interior (which is, however, linearizable). The effects of the rotation are also easily measurable and this is one of the areas in which helioseismology has contributed very significantly to our understanding of the solar interior.

Most inversions of both real and artificial data have been using so-called 1- or 1.5-dimensional inversion methods. When using the 1-dimensional methods it is assumed that the rotation rate is independent of latitude (or the be more accurate, only a particular weighted integral of the rotation rate, as a function of latitude, is obtained at each radius). While this is manifestly wrong (as is obvious from the surface rotation rate being $\approx 50\%$ larger at the equator than at the pole), this problem is particularly simple and thus makes it a good tool for understanding various properties of the inversions. Also the 1-dimensional inversion problem forms the basis of the 1.5-dimensional problem, which has been used extensively for analyzing observations.

In the 1.5-dimensional inversion problem (described in some detail in Schou, Christensen-Dalsgaard, and Thompson, 1993, cf. App. 3), the latitude dependence of rotation as well as the m dependence of the frequencies are expanded in polynomials, yielding a series of 1-dimensional inverse problems.

The full 2-dimensional problem, on the other hand, uses a general 2-dimensional inversion law and the frequencies of the individual modes. This problem is considerably more complicated and is at the moment only at an experimental stage.

In the following I will start with a discussion of the 1-dimensional inversions, as they are considerably simpler than the 1.5-dimensional or 2-dimensional inversions. Since the 1.5-dimensional and 2-dimensional problems are both covered in significant detail in Schou, Christensen-Dalsgaard, and Thompson (1993, cf. App 3), I will not go into too many details here.

4.1 1-Dimensional Inversions

The 1-dimensional problem has been covered in a substantial number of papers, in particular the paper by Christensen-Dalsgaard, Schou, and Thompson (1990, cf. App. 1) (in the following CDST) compared several of the popular methods.

As already mentioned, the 1-dimensional methods are based on the assumption that the solar rotation is independent of latitude, and that the inverse problem is therefore reduced to a purely radial problem. This version of the forward problem was mentioned briefly in the section (2.3) on the effects of rotation on solar oscillations. Examples of mode kernels for this problem have been shown in figure 4.1. Notice that the kernels are dominated by the square of the radial displacement, cf. the chapter on solar oscillations.

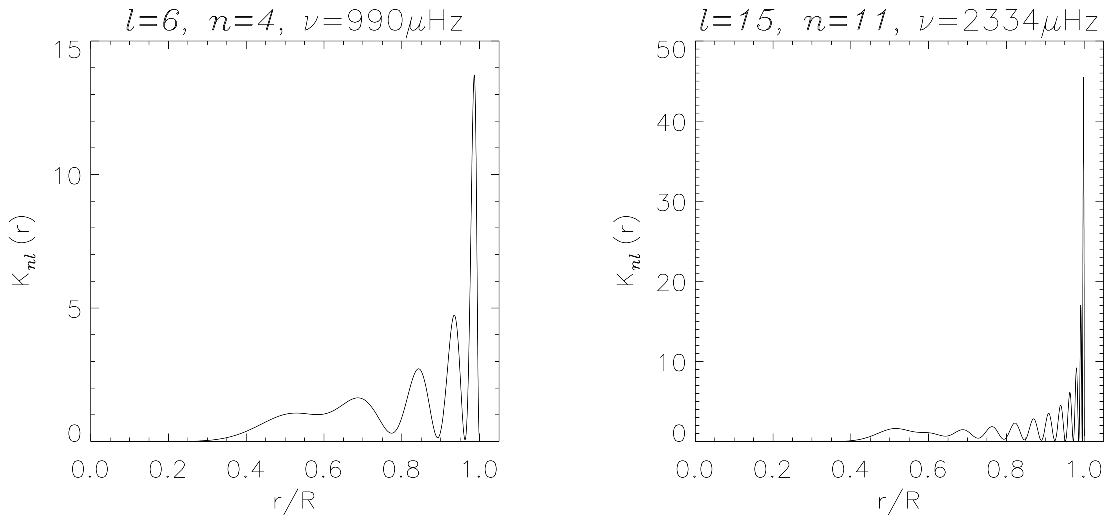

Figure 4.1: Mode kernels for selected modes. The radial displacements for these modes are shown in chapter 2.

In order to simplify the problem the modes will be labeled with a single subscript $_i$ instead of a double subscript $_{nl}$. With these simplifications the forward problem can be written as

$$\omega_i = \int K_i(r)\Omega(r)dr, \tag{4.1}$$

where Ω is the rotation rate, $K_i(r)$ are kernels, which can be found from (2.21) given the mode eigenfunctions, and ω_i are the observed splittings. The observed splittings are obviously not error free and an estimate of the covariance matrix E_ω is provided. Typically the errors on the individual splittings are assumed independent such that E_ω is a diagonal matrix with $(E_\omega)_{ii} = \sigma_i^2$ where σ_i are the assumed standard errors on the individual splittings.

The inverse problem then consists of finding an estimate $\bar{\Omega}$ of the true solar rotation rate Ω, taking into account the observed splittings and their associated errors. Since only a finite set of observations exists and Ω is a general (presumably?) continuous function, it is obviously not possible to find Ω, even in the absence of observational errors. Hence the problem is ill-posed and it is necessary to impose other constraints on the inferred rotation rate. The different inversion methods impose different constraints; in some methods the constraints are explicitly imposed, in others they are more indirectly present. Typically, the constraints are on the smoothness of the solution.

Since the formulations for the different inversion methods are very different, it is useful to have a set of diagnostics telling, in a uniform way, how the observations are transformed into the result for each of the methods. It turns out that for the class of linear inversion methods (which includes most methods currently used) a number of such diagnostics exists. A linear inversion method is defined as one in which there exists, for each r_0, coefficients $c_i(r_0)$ such that

$$\bar{\Omega}(r_0) = \sum_i c_i(r_0)\omega_i \tag{4.2}$$

independent of the observations ω_i, but not of the assumed errors, parameters used in the inversion, etc.

One diagnostic is obviously the coefficients telling which modes contributed to the inferred rotation rate at any given point.

Substituting eq. (4.1) (without adding the observational errors) into eq. (4.2) one obtains

$$\bar{\Omega}(r_0) = \sum_i c_i(r_0)\omega_i = \sum_i c_i(r_0) \int K_i(r)\Omega(r)dr = \int K(r_0, r)\Omega(r)dr, \tag{4.3}$$

where

$$K(r_0, r) = \sum_i c_i(r_0)K_i(r) \tag{4.4}$$

are the so-called averaging kernels, describing how the underlying rotation rate is transformed into the inferred rotation rate in the absence of observational errors. Examples of averaging kernels have been shown in figure 4.2.

Given the averaging kernels it is possible to define several other useful measures. Some of the most useful measures are the area under each averaging kernel

$$uni(r_0) = \int K(r_0, r)dr \tag{4.5}$$

and the centre of gravity

$$\bar{r}(r_0) = \frac{\int rK(r_0, r)dr}{uni(r_0)}. \tag{4.6}$$

The usefulness of these measures comes from the fact that if $uni = 1$ for all r_0 then a constant rotation rate will be reproduced in the absence of observational errors. Most inversion methods are designed to give the averaging kernels unit area. Furthermore if $r_0 = \bar{r}$ or the inferred rotation rate is taken as a function of $\bar{r}$ instead of r_0, a rotation rate which is a linear function of radius will be reproduced exactly.

In general it is useful to define the so-called moments

$$m_j(r_0) = \frac{1}{j!} \int (r - r_0)^j K(r_0, r)dr. \tag{4.7}$$

The usefulness of these moments comes from the fact that if the true rotation rate Ω is expanded in a Taylor series around a particular r_0

$$\Omega(r) = \Omega(r_0) + (r - r_0)\frac{d\Omega}{dr}(r_0) + \frac{1}{2!}(r - r_0)^2 \frac{d^2\Omega}{dr^2}(r_0) + \ldots \tag{4.8}$$

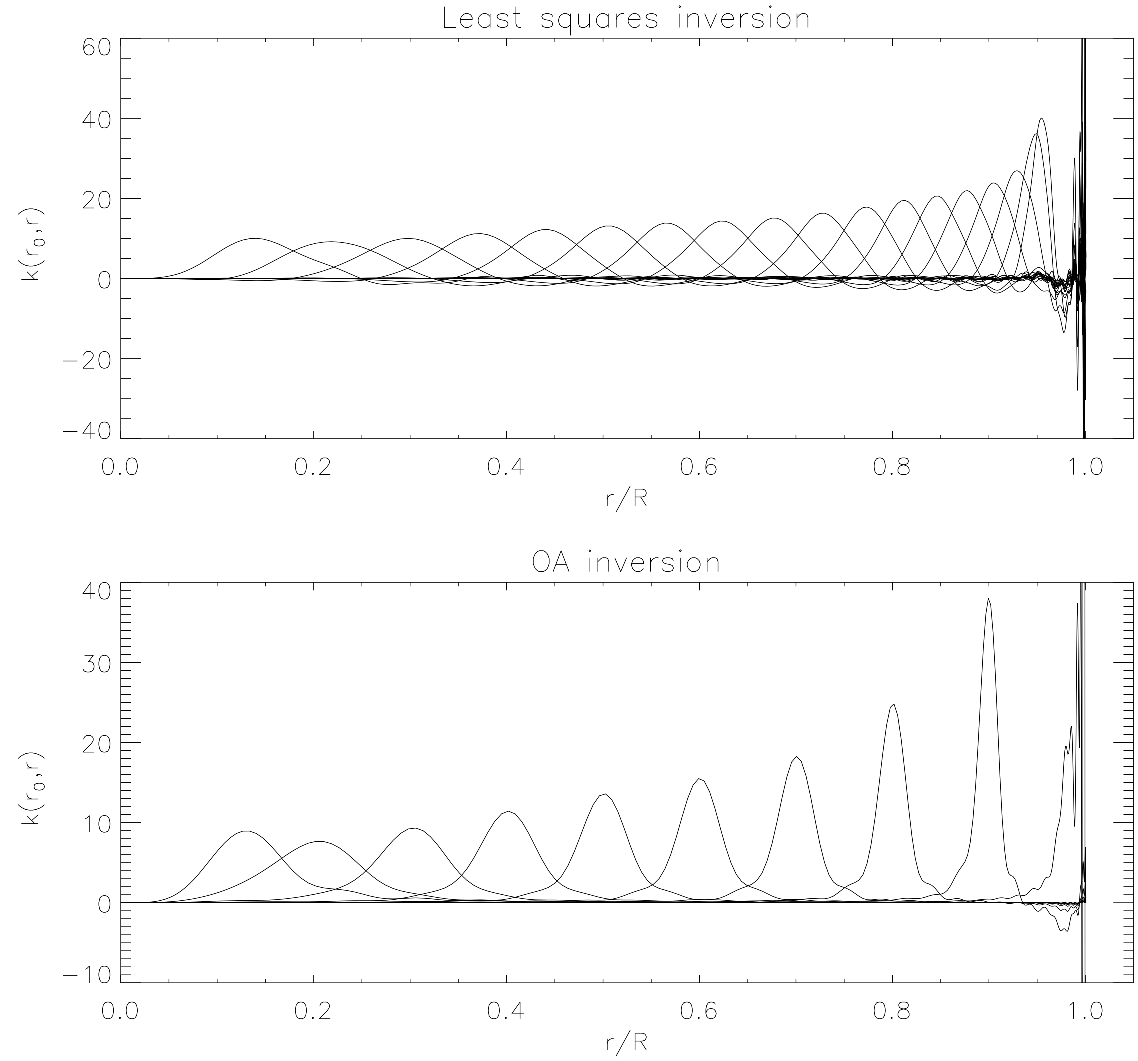

Figure 4.2: Examples of averaging kernels for a least squares inversion of set1 from CDST (which contained modes with $1 \leq l \leq 200$) with $\mu = 10^{-6}$ (top plot), using second derivative smoothing, and for an optimally localized averaging kernels method with $\mu = 10^{-2}$ (bottom plot).

the inferred rotation rate can be expressed as

$$\begin{aligned}\bar{\Omega}(r_0) &= \int K(r_0,r)\left(\Omega(r_0)+(r-r_0)\frac{d\Omega}{dr}(r_0)+\frac{1}{2!}(r-r_0)^2\frac{d^2\Omega}{dr^2}(r_0)+...\right)dr \\ &= m_0(r_0)\Omega(r_0)+m_1(r_0)\frac{d\Omega}{dr}(r_0)+m_2(r_0)\frac{d^2\Omega}{dr^2}(r_0)+...\end{aligned} \tag{4.9}$$

Hence if Ω is well represented by a Taylor series, $m_0 = 1$ and $m_j \approx 0$ for $j > 0$, then $\bar{\Omega}$ will be a good approximation to Ω. Note that $uni = m_0$ and that $\bar{r} = r_0$ is equivalent to $m_1 = 0$. If Ω is *not* well represented by a Taylor series within the area where $K(r_0,r)$ is significantly nonzero, this argument breaks down and $\bar{\Omega}$ is not necessarily a good approximation to Ω when the moments are small.

Notice that while it is possible to force m_0 and m_1 to 1 and 0 respectively by renormalizing the kernels (and the corresponding $\bar{\Omega}(r_0)$) and redefine r_0 to $\bar{r}$, it is not possible to also force $m_2 = 0$. Also note that $m_2 = 0$ means that the averaging kernels must be negative somewhere, typically in the form of negative sidelobes, and hence will introduce ringing in the inferred rotation rate at near discontinuities in the real rotation rate.

A useful measure for the width of the averaging kernel is the so-called spread of the kernel given by

$$\mathrm{spread}(r_0) = 12 \int K(r_0, r)^2 (r - r_0)^2 dr \tag{4.10}$$

(see the description of the optimally localized averages method in Christensen-Dalsgaard, Schou and Thompson, 1990 (cf. App. 1) for another use of this measure), the reason for the peculiar normalization being that a 'boxcar' function with a width a (and height $1/a$ to make it unimodular) will have a spread of a if r_0 is the centre of the boxcar. This measure of the width puts a fairly heavy penalty on the averaging kernels having any structure far from r_0.

Another thing one can calculate from the coefficients is the covariance matrix E of the different points on the inferred rotation curve

$$E(r_0, r_0') = \sum_i \sigma_i^2 c_i(r_0) c_i(r_0') \tag{4.11}$$

where it has been assumed that the covariance matrix for the observations is diagonal. From this the error on the inferred rotation rate at a point r_0 can be found from

$$\sigma(\bar{\Omega}(r_0)) = \sqrt{E(r_0, r_0)} = \sqrt{\sum_i \sigma_i^2 c_i(r_0)^2}. \tag{4.12}$$

In case of identical errors σ on the original data points, which has often been used for artificial data, the so-called error magnification Λ is defined by

$$\sigma(\bar{\Omega}(r_0)) = \Lambda(r_0)\sigma, \tag{4.13}$$

that is

$$\Lambda(r_0) = \sqrt{\sum c_i(r_0)^2} \tag{4.14}$$

is the ratio of the errors in the output rotation curve to the input errors.

The 'normalized' covariance matrix given by

$$E'(r_0, r_0') = \frac{E(r_0, r_0')}{\sigma(\Omega(r_0))\sigma(\Omega(r_0'))} \tag{4.15}$$

is often more useful as the off-diagonal elements tell how strongly correlated two points in the inferred rotation rate are.

For nonlinear inversion methods similar diagnostics are harder to define. Some of the diagnostics can be defined differentially by linearizing the solution, but the results are obviously more difficult to interpret. For all methods it is desirable to at least have estimated errors on the inferred rotation rate, but these are often provided by the nonlinear methods and can in most cases be found using a Monte Carlo simulation.

The different inversion methods generally work by making a trade-off between two quantities, such as good resolution/small errors on the inferred rotation rate, or fitting the data/limiting oscillations of the inferred rotation rate. For this and other reasons it is often informative to plot these two quantities against each other, varying the parameter (generally called a trade-off parameter) describing the trade-off. Examples of such trade-off curves have been shown in figure 4.3. While one type of trade-off curve typically corresponds to one particular inversion method, it is sometimes useful to plot the same quantities, or even quantities not used in a particular inversion method, for different methods (see eg. CDST).

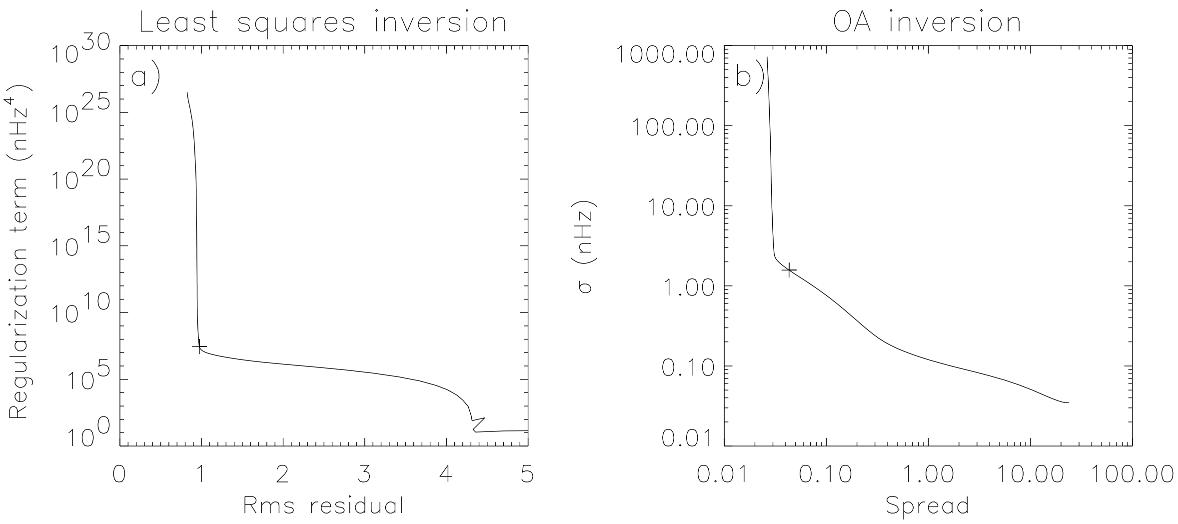

Figure 4.3: Examples of trade-off curves for the mode set used for figure 4.2. a) is for a regularized linear least squares inversion using second derivative smoothing, b) for the optimal averaging kernels method at $r_0 = 0.5R$. The underlying rotation rate used had $\Omega = 400nHz$ for $r/R < 0.65$, varied linearly between $400nHz$ and $500nHz$ for $0.65 \leq r/R \leq 0.75$ and had $\Omega = 500nHz$ for $r/R > 0.75$. The splittings had normally distributed errors added. The points corresponding to the inversions shown in figure 4.2 are indicated by crosses. The jumps in the curve for the least squares method where the rms > 4 are due to numerical problems. For both plots units such that $R \equiv 1$ (dimensionless) have been used

Figure 4.3a shows a trade-off curve for the regularized least squares method. The behaviour is the classic one for this method. For high μ, the solution is totally dominated by the regularization term, forcing the solution to be a straight line (very small regularization term). As μ is decreased (going left), the solution is allowed to match the data better, causing a rapid decrease in the residual. Eventually (at rms ≈ 1) the residuals become dominated by the random errors, and they become almost impossible the reduce, even by allowing the solution to oscillate wildly (causing a large regularization term). If it were not for numerical problems, it would be possible to reduce the rms to zero (forgetting about accidental degeneracies of the mode kernels) by allowing a sufficiently rapidly varying solution. Generally the regularization parameter is chosen close to the knee in the trade-off curve around rms=1, as indicated by the cross.

Figure 4.3b shows a trade-off curve for the optimally localized averaging kernels method. Note that as this is a 'local' method in that the trade-off is done locally, only the results for a single r_0 have been shown. Again the trade-off curve is classic. For high μ (to the right on the plot), the averaging kernels are very broad, but the random error on the inferred rotation at the target point is very low. As the trade-off parameter is reduced, the kernels become narrower, at the expense of an increased random error. For a sufficiently small trade-off parameter, the width of the averaging kernels can no longer be reduced; the resolution limit for the modeset has been reached. Reducing μ further only increase the error without improving the resolution substantially. Again the trade-off parameter is generally chosen around the knee in the curve.

4.2 1.5-Dimensional Inversions

As already mentioned, the Sun is not rotating as a solid body. Rather there are large gradients in the rotation rate as a function of both radius and latitude. Fortunately, the problem is still linear in that it is possible to find kernels $K_{nlm}(r,\theta)$ such that the perturbation $\Delta\omega_{nlm}$ in the mode frequency is given by

$$\Delta\omega_{nlm} = \int_0^R \int_0^\pi K_{nlm}(r,\theta)\Omega(r,\theta)drd\theta, \tag{4.16}$$

Generally, the observed frequencies have not been given as individual mode frequencies, but rather as an 'unperturbed' mode frequency ω_{nl} and a series of so-called a-coefficients in a fit to the frequencies of the form

$$\omega_{nlm} = \omega_{nl} + 2\pi \sum_{j=1}^{j_{max}} a_j(n,l)\,\mathcal{P}_j^{(l)}(m)\ , \tag{4.17}$$

where $\mathcal{P}_j^{(l)}$ is a polynomial of degree j. These polynomials have traditionally (in all previously published observations) been Legendre polynomials $\mathcal{P}_j^{(m)} = lP_j(m/l)$ (or in some cases using L instead of l) but they may also be chosen as in Ritzwoller and Lavely (1991). For more details on how these polynomials may be chosen and calculated, see Schou, Christensen-Dalsgaard, and Thompson (1993, cf. App. 3).

Having expanded the splittings in polynomials it is natural also to expand the rotation rate in polynomials in x^2, where $x = \cos\theta$:

$$\Omega(r,\theta) = \sum_{s=0}^{s_{max}} \Omega_s(r)\psi_{2s}(x), \tag{4.18}$$

where ψ_{2s} is a polynomial in x of degree $2s$. As mentioned in Schou, Christensen-Dalsgaard, and Thompson (1993, cf. App. 3) it may then be shown that kernels K_{nls}^j exist, such that

$$2\pi a_{2j+1}(n,l) = \sum_{s=j}^{s_{max}} \int K_{nls}^j(r)\Omega_s(r)dr; \tag{4.19}$$

that is, only Ω_s's with $s \geq j$ contributes to a_{2j+1}. This equation forms the basis for the 1.5-dimensional inversions.

Rewriting the previous equation gives

$$\int K^j_{nlj}(r)\Omega_j(r)dr = 2\pi\tilde{a}_{2j+1}(n,l) = 2\pi a_{2j+1}(n,l) - \sum_{s=j+1}^{s_{max}} \int K^j_{nls}(r)\Omega_s(r)dr \qquad (4.20)$$

making it possible to do the full inversion by starting with the inversion for the term with $j = s_{max}$ and working downwards using the previous inversions to subtract (an approximation to) the sum on the right-hand side. Clearly, it is less than ideal to use the (smoothed and error contaminated) inferred rotation rates for the higher order terms for this purpose, but as the terms are generally very small (at least for $l \gg 1$), this is only a minor problem.

As shown in Ritzwoller and Lavely (1991), the second term on the right-hand side of the previous equations vanishes if the polynomials they suggest (or ones similar to theirs, as in Schou, Christensen-Dalsgaard, and Thompson 1993, cf. App. 3) are used. This is the justification for using their polynomials or polynomials proportional to theirs instead of the traditional ones, in eg., the analysis of the Fourier Tachometer data.

As this method is just a series of 1-dimensional inversions, it is for each term possible to define coefficients, averaging kernels etc. As discussed in Schou, Christensen-Dalsgaard, and Thompson (1993, cf. App. 3), many of the concepts discussed in connection with the 2-dimensional inversions are also relevant here, as this is after all, essentially a 2-dimensional method.

4.3 2-Dimensional Inversions

In the full 2-dimensional inversions neither the rotation rate nor the splittings are expanded in polynomials. Instead a general rotation rate and the raw frequency splittings are used (although it is possible to use a-coefficients).

Computationally there are two basic problems facing 2-dimensional inversions. One is that the use of all the individual splittings instead of just a small number of a-coefficients means that the number of observations is larger by a factor of the order l. For the mode set considered in Schou, Christensen-Dalsgaard, and Thompson (1993, cf. App. 3), the number of (nl)s was 2413, while the number of (nlm)s was 195870!

The other basic problem is that the number of gridpoints needed to represent a full 2-dimensional rotation rate with the same resolution as in the 1-dimensional case is again close to two orders of magnitude larger than in the 1-dimensional case.

In theory most of the methods considered for the 1-dimensional inversions could be extended to the 2-dimensional case. However, given the increased computational burden just mentioned, some of the methods become essentially impossible. In particular, the optimal averaging kernels method, which is one of the favorite methods in the 1-dimensional case, becomes very difficult, as it requires solving linear systems of the order the number of splittings (≈ 200000 in the case just mentioned).

Of the 1-dimensional methods left, both the regularized least squares method (Schou, 1991b, Schou, Christensen-Dalsgaard, and Thompson 1993 (cf. App. 3), Sekii 1990, 1991) and an asymptotic method (Sekii 1993) have been successfully implemented in 2-dimensional versions. In the method developed by Sekii, the inversion for the latitude dependence is done first using all ms for each (nl), followed by an inversion for the radial dependence for each target latitude.

As in the 1-dimensional method, it is possible to devise various measures of resolution, etc., some of which have been discussed in Schou, Christensen-Dalsgaard, and Thompson (1993, cf. App. 3).

It should be mentioned that there are several promising ways to improve on the computational problems. For the least squares method, experience from the 1-dimensional case indicates that the inversions at points separated by much more than the resolution are largely independent. If this result also holds in the 2-dimensional case, it would allow one to perform the full inversion by doing a number of inversions for different areas of the $r-\theta$ plane.

Another possibility is to combine the splittings before the inversion, such as to reduce the number of input data points to the inversion(s). Two approaches that have been considered are to combine splittings given the information from the 1-dimensional latitude inversions, as discussed in Schou, Christensen-Dalsgaard and Thompson (1993, cf. App. 3). A special case of this approach is that used by Sekii (1993) discussed above and in App. 3.

It should also be possible to use the ideas using (G)SVD methods presented in Christensen-Dalsgaard, Hansen, and Thompson (1992) to combine splittings in advance. In this method the fact that the mode kernels are highly degenerate is used to combine the splittings in such a way as to preserve the useful information in the splittings and rejecting those combinations that are ill determined.