

INVERSION METHODS

M.J. THOMPSON

*Astronomy Unit, School of Mathematical Sciences,
Queen Mary & Westfield College, London, E1 4NS, England*

1. Introduction

I want to start by addressing the question, ‘What is inversion?’ My answer would be that inversion is the process of going from data to making inferences about the object under study. In the case of helioseismology, the data at the present time are principally the mode frequencies, and the object under study is the solar interior.

For example, suppose we found that observed frequencies of p modes confined to the convection zone agreed well with the frequencies of a theoretical model, but that the observed frequencies of modes penetrating just beneath the convection zone were higher than those in the model. We might reasonably infer that there was something wrong about our model just beneath the base of the convection zone. Moreover, since for acoustic modes the frequency $\nu \sim A/\int c^{-1}ds$, where c is the sound speed, the integral is with respect to distance s along a ray path, and A is just some constant of proportionality, we might infer that the sound speed was too small in this region in the model. Since we know that increasing the opacity in the model in this region would have the effect of increasing the sound speed there, we might infer therefore that the opacity in the model might need to be increased in this region.

In my view, this hypothetical scenario (based though on the actual findings of Christensen-Dalsgaard et al. 1985) is already an example of inversion. It illustrates how knowledge about the details of the mode physics – the depth of penetration of different modes, and the role of the sound speed – is used to make localized inferences about possible errors in the model. In a real case, one would also make some quantitative estimate of the shortfall in the model sound speed, and give proper consideration too to the uncertainties on that inference.

The frequencies of the Sun’s global modes depend on conditions in the interior through adiabatic exponent Γ_1 , pressure p and density ρ – all functions of position \mathbf{r} . The observed p modes depend on these three quantities principally in the combination $c^2 = \Gamma_1 p/\rho$. They also depend on the interior velocity field $\mathbf{v}(\mathbf{r})$ (in particular the rotation) and magnetic field $\mathbf{B}(\mathbf{r})$. These – and functions that can be derived just from a knowledge of these quantities – are the primary seismic variables about which we can make helioseismic inferences. Inferences about secondary quantities, such as the opacity in the example above, which affect the modes only through influencing the primary seismic variables, are at a different level, since there may be other physical causes that may change lead to the same change in the primary variables.

The simple example above also illustrates that it is necessary to have a framework for solving the forward problem in order to perform an inverse analysis. Only if one can predict (albeit perhaps in some approximation) what data would be produced by a particular physical model can one hope to use the data to make inferences about that physics.

The dependence of the mode frequencies on the structure and dynamics of the solar interior is intrinsically nonlinear. Most helioseismic inversion is based on some form of linearization, and for that reason I shall concentrate on the linearized problem for most of this review. However, I should at least note that there are other approaches. One such might be described as inversion by doing the forward problem, namely changing some aspect of the model, seeing what effect this has on the frequencies and thus finding which model amongst some set best fits the observational data. This approach is known as model calibration. Recently there have been several attempts at what is in effect an implementation of this approach with an automated search through some set (or space) of models. Specifically Paul Charbonneau at HAO and Barry Laphorn, a student of mine, have both been looking at using genetic algorithms for inverting for the solar internal rotation (see, e.g. Charbonneau et al. 1997). The approach looks promising, at least when only a few dozen independent parameters are used to parametrise the space of models to be explored. Another set of nonlinear methods is based on using an asymptotic description of the forward problem, usually leading to having to invert an Abel integral (e.g. Gough 1984, 1986; Brodsky & Vorontsov 1988; Shibahashi 1988).

In the context of inversions for radial hydrostatic structure, linearization is generally performed by formally perturbing a variational principle for the frequencies. Retaining only first-order perturbation quantities, one obtains a linear functional equation that relates differences in global frequencies between the Sun and a model to differences in their primary variables (e.g. Gough 1985). Further assumptions and manipulations can be used to relate the frequency differences instead to differences in secondary quantities such as opacity, chemical abundances or temperature.

It is rather natural also to handle the effects on the mode frequencies of the internal magnetic and velocity fields, and in particular the internal rotation, in a linearized approximation. For example, the rotation period of roughly one month is very much longer than the typical period of a mode, which is around five minutes, so that including the effects of rotation on the mode as a linearized perturbation is a very good approximation.

As a very simple example of a set of linear constraints to be inverted, I introduce the archetypal equation

$$d_i = \int_0^1 K_i(x)\Omega(x) dx + \epsilon_i \quad (i = 1, 2, \dots, M). \quad (1)$$

Here d_i represent the M pieces of observational data, $\Omega(x)$ the unknown function about which we wish to make inferences, ϵ_i the errors in the data, and $K_i(x)$ are known functions. For simplicity I take the errors to be independent normally distributed with zero mean and with variance σ_i^2 . The linearized problem for inferring the Sun's internal rotation is very similar to this, except that the function Ω depends on radius and latitude rather than on just one variable x : the data are the so-called frequency splittings between modes with different azimuthal degrees m . The linearized problem for inferring the differences in radial structure between Sun and reference model is also of this form, except that the structural differences generally need to be represented

by more than one unknown scalar function $\Omega(x)$: the data are then differences (or relative differences) between observed frequencies and those of the reference model.

2. Data fitting

An obvious way to solve the observational constraints represented by equation (1) is to use the least-squares method. Thus one might seek a representation of Ω of the form

$$\bar{\Omega}(x) = \sum_{j=1}^N \bar{\Omega}_j \phi_j(x)$$

where $\phi_j(x)$ ($j = 1, \dots, N$) are chosen basis functions, and the constants $\bar{\Omega}_j$ are determined by minimizing the data mismatch

$$\sum_{i=1}^M \left[d_i - \int_0^1 K_i \bar{\Omega} dx \right]^2 / \sigma_i^2. \quad (2)$$

In practice this is not very successful unless N is chosen to be very small. The reason is that the data depend on the unknown Ω through $\int K_i \Omega dx$: this tends to smooth variations in Ω , just as straightforward integration of a function tends to produce a result that is smoother than the original function itself. Inversion – inferring Ω from the data d_i – is analogous to the inverse of integration, namely differentiation; and as is well known, differentiating noisy data means that the noise tends to get magnified. The same happens with a naïve application of least squares to the helioseismic inversion problem: if the base functions permit it, the solution will generally be dominated by rapidly varying oscillations caused by the solution attempting to fit the noise in the data.

A commonly employed resolution of this problem is to introduce some regularization into the least-squares method. Instead of minimizing expression (2) one instead minimizes an expression such as

$$\sum_{i=1}^M \left[d_i - \int_0^1 K_i \bar{\Omega} dx \right]^2 / \sigma_i^2 + \lambda^2 \int_0^1 (\bar{\Omega}'')^2 dx. \quad (3)$$

Here the double prime denotes second derivative with respect to x , and the inclusion of the new term penalizes solutions which have large second derivatives. The value of the constant λ is chosen so as to trade-off the essentially opposing aims of minimizing the data mismatch and minimizing the penalty term. The choice of the exact form of the penalty term is rather ad hoc, but the advantage of making it quadratic in $\bar{\Omega}$ is that one is still led to a set of linear equations to be solved to find the coefficients $\bar{\Omega}_j$.

It is illuminating to see how a Bayesian would analyse the difficulty we encounter with unregularized least-squares and our ad hoc solution of regularized least squares. Bayes's Theorem says that

$$P(\Omega|DX) \propto P(D|\Omega X)P(\Omega|X),$$

The expression on the left is the *a posteriori* probability of Ω , given the data D and any prior information (or prejudice) X . The first term on the right is the probability

of observing the data D given Ω and any prior information X , and the remaining term is the prior probability of Ω before the data is obtained. Given our assumption about the data errors being independent and normal,

$$P(D|\Omega X) \propto \prod_{i=1}^M e^{-(d_i - \int K_i \Omega dx)^2 / 2\sigma^2},$$

so $-2 \ln P(D|\Omega X)$ is precisely the expression (2). Thus minimizing (2) is equivalent to maximizing the likelihood of the data. If we accept that what we should be doing in fact is finding the function Ω that maximizes $P(\Omega|DX)$, then we see that unregularized least squares is equivalent to assuming that the prior probability $P(\Omega|X)$ is the same for all Ω . Maximizing $P(\Omega|DX)$ is equivalent to minimizing $-2 \ln P(\Omega|DX)$ which, apart from an additive constant, is equal to the usual least-squares term plus $-2 \ln P(\Omega|X)$. By equating this with expression (3), it is evident that the regularized least squares is equivalent to assuming a prior probability satisfying $-2 \ln P(\Omega|X) = \lambda^2 \int (\Omega'')^2 dx$ plus constant, i.e.

$$P(\Omega|X) \propto e^{-\frac{1}{2} \lambda^2 \int (\Omega'')^2 dx}.$$

This then is our prior probability for Ω . It embodies a prejudice that functions Ω with large second derivatives are less probable (prior to getting any data to the contrary). However, its form has no firm justification. Could we instead introduce a prior which had physical motivation, perhaps updated in the light of previous helioseismic inference?

3. Characterizing all functions that are consistent with the data

The regularized least-squares method at best provides just one example of a function that is consistent with the data. (I say ‘at best’ because if the value of the chi-squared per degree of freedom is not close to unity then the method fails even to do that.) Because the data provide only a finite number of linear constraints (1), there must be an infinite number of other solutions that are equally consistent with the data.

An alternative approach to the inversion problem is to try to characterize all functions that are consistent with the data. It is almost tautological to say that the set of original constraints (1) provides such a characterization. And each of the original constraints individually provides a partial characterization, since all solutions must satisfy each such constraint. It follows then that any linear combination of the original constraints also provides a partial characterization: any function consistent with the original data must satisfy

$$\sum_{i=1}^M c_i d_i = \int_0^1 \mathcal{K} \Omega dx + \sum_{i=1}^M c_i \epsilon_i, \tag{4}$$

where

$$\mathcal{K} = \sum_{i=1}^M c_i K_i(x), \tag{5}$$

for any constants c_i we choose. Why is this progress? We have still just got an integral constraint on Ω . But in helioseismology the original kernels K_i are oscillatory functions

with substantial amplitude over perhaps a large part of the solar interior. Suppose that with the freedom to choose the constants c_i one could arrange to make the averaging kernel \mathcal{K} localized close to some point $x = x_0$ and moreover nicely peaked there, with unit area when integrated from $x = 0$ to $x = 1$. Then the localized average of any solution Ω , weighted by the averaging kernel \mathcal{K} , would have to satisfy constraint (4), the left-hand side of which is determined by the data. It turns out that in the various helioseismic applications and over a large part of the Sun it is indeed possible to construct beautifully localized averaging kernels. Each of these corresponds to a constraint of the form (4) on the set of all possible solutions, which I at least find easier to comprehend than the original data constraints. Each of the constraints says that although there may be an infinite set of functions Ω consistent with the data, in the vicinity of x_0 where one has succeeded in constructing a well-localized averaging kernel any solution Ω must have a local average (made precise by the weighting \mathcal{K}) equal to the left-hand side of (4), modulo the error term $\sum c_i \epsilon_i$ whose statistical properties can be established if the statistics of the original data errors ϵ_i are known.

There are various ways in which the coefficients c_i can be found in order to construct averaging kernels with desired degree of localization. Such optimally localized averages methods are based on ideas by Backus & Gilbert (1968), and are described in the helioseismic context by e.g. Gough (1985) and Pijpers & Thompson (1992). As with regularized least squares, there is a trade-off to be made, since trying to localize the averaging kernel too much would increase the variance of the error term $\sum c_i \epsilon_i$ and would eventually render the constraint (4) practically useless.

4. Resolution

If it is possible to localize an averaging kernel in the vicinity of x_0 , with a substantial positive peak there and with little amplitude elsewhere, then one can define the resolution in the x -direction attained by the inversion in terms of some measure of the width of the averaging kernel peak, e.g. its full width at half maximum (FWHM). In the case of a 2-D inversion, this can be done in both the radial and latitudinal directions.

With a quadratic penalty term, as in eq. (3), the regularized least-squares (RLS) method is linear, in the sense that the solution at any point x_0 is a linear combination of the data d_i . Thus implicitly coefficients $c_i(x_0)$ exist, and so the RLS solution at x_0 is actually also an average of the true function Ω weighted by an averaging kernel of the form (5). Likewise, the solution from any linear data-fitting method can be interpreted point-by-point in terms of averaging kernels and linear functionals characterising the set of all functions that are consistent with the data. Of course the RLS averaging kernels may not be so nicely localized – after all, one did not explicitly seek to localize them. However, in practice they often are, albeit perhaps with negative sidelobes adjacent to the main positive peak, and so for the RLS too one can define attained resolution in terms of the width of the averaging kernel's main peak in regions where the kernel is well-localized. And if over a range of different target locations x_0 it is possible using the explicit averaging kernel construction methods to produce well-localized kernels, why not interpret the solutions $\sum c_i(x_0)d_i$ as forming a continuous curve – a solution – which probably does not fit the data but which can rather be interpreted as a smoothed-out approximation to the true $\Omega(x_0)$.

The standard deviation of the error $\sum c_i(x_0)\epsilon_i$ in the solution is $(\sum c_i^2 \sigma_i^2)^{1/2}$. The coefficients c_i depend on any adjustable parameters in the method, and there-

fore so too do the error standard deviation and the attained resolution as defined above. One can then define at a given location x_0 a trade-off curve of error standard deviation against resolution, for any given method. As one tries to squeeze the coefficients to get better resolution, the error increases; and, conversely, as one tries to reduce the error, the resolution gets poorer. Not only do the trade-off curves for different linear helioseismic inversion methods have the same qualitative behaviour, Christensen-Dalsgaard et al. (1990) found that they were very similar quantitatively, in regions where well-localized kernels could be constructed. The resolution, in the sense defined above, cannot be reduced to arbitrarily fine values. In fact there is a fundamental limit on how fine the resolution can be, which is determined by the shortest scales of variation of the kernels K_i in the vicinity of x_0 . For a primary variable such as rotation or sound-speed, this scale is directly related to local wavenumber of the mode eigenfunctions. Within the propagation region of p modes the wavenumber is related to c/ω , where c is adiabatic sound speed evaluated at location x_0 and ω is the angular frequency of the mode. Thus the finest attainable resolution at x_0 when inverting p-mode data for rotation or sound speed is c/ν_{\max} times a number of order unity (to allow for the conversion from eigenfunction to kernel), ν_{\max} being the top limit on the range of frequencies in the data set. The practical resolution limit may be poorer than this because of the noise in the data, which prevents the averaging kernel from being so well localized while keeping the error standard deviation within acceptable bounds.

The best resolution in the convection zone is probably now achieved with the SOI-MDI dataset. Just to give some idea of the currently attainable resolution, I quote some values for the typical resolution achieved with the 144-day MDI set of splitting coefficients in the region where averaging kernels can be reasonably well localized (values are for the RLS inversion of Schou et al. 1997a; but values for 2-D optimally localized averages would be similar). At 30° latitude, for example, the FWHM radial widths of the averaging kernels (i.e. radial resolution) are $0.02R$ at $r = 0.965R$ (R being the solar radius), $0.05R$ at $r = 0.82R$, $0.08R$ at $r = 0.72R$ and $0.11R$ at $r = 0.52R$; the corresponding FWHM widths measured in the latitudinal direction (i.e. latitudinal resolution) are $0.13R$, $0.14R$, $0.20R$ and $0.29R$. These results are summarised in Table 1. Widths vary less with latitude than with radius: at radius $0.72R$, for example, the FWHM radial widths at different latitudes are $0.07R$ at 15° latitude, $0.08R$ at 45° and $0.09R$ at 60° . The latitudinal widths hardly vary with latitude at fixed radius, in the region where the averaging kernels are reasonably well localized. The latitudinal resolution is poorer than the radial resolution, probably because the splittings are only available as so-called a-coefficients rather than individual splittings. Nonetheless, with a-coefficients up to a_{36} (see Schou et al. 1997b), the latitudinal resolution is the best yet achieved.

Table 1. Radial and latitudinal resolution attained with a 2-D RLS inversion of the 144-day MDI set of splitting coefficients, as measured by averaging kernel widths at selected target locations.

Target radius	Target latitude	Radial width(FWHM)	Latitudinal width(FWHM)
$0.965R$	30°	$0.02R$	$0.13R$
$0.82R$	30°	$0.05R$	$0.14R$
$0.72R$	30°	$0.08R$	$0.20R$
$0.52R$	30°	$0.11R$	$0.29R$

Note that if one quotes interquartile widths instead of FWHM (e.g. Christensen-Dalsgaard et al. 1990) the interquartile widths are smaller by a factor of about 0.57 when the shapes of the kernels are approximately Gaussian.

Before leaving the subject of resolution, I would like to point out that what one might call super-resolution can be achieved, but only by the imposition of additional assumptions in the inversion. If for example in a hypothetical problem the underlying function Ω were assumed to be a delta function, and if the averaging kernels were (for simplicity) nice Gaussians whose properties varied only slowly with target location x_0 , then the position of the delta function could be inferred from where the solution (4) – viz the convolution of the kernels with the delta function – attained its maximum. Provided the data noise were sufficiently small, the position of this maximum could be located much more precisely than one FWHM of the Gaussian kernels. I call this super-resolution. But features have not really been resolved beneath the lengthscale given by the width of the kernels. If for example, the true Ω had in fact contained two closely-spaced delta functions, these would not have been resolved. A similar example, of more relevance to our applications, is if the true Ω were assumed to have some simple form locally like a step function, then the position at which the step occurs could be localized more precisely than one kernel width. Helioseismic examples where specific forms have been assumed include the attempts to locate the base of the convection zone (Christensen-Dalsgaard et al. 1991) and to measure the position and thickness of the tachocline (e.g. Kosovichev 1996). In these cases, a particular form of the radial structure or rotation profile were assumed, and the data then allowed the location of the assumed form to be determined rather precisely. However, these regions might have a more complex structure not represented by the simple assumed form and unresolved by the inversions.

5. On interpreting results of inversions

When presented with the result of a linear helioseismic inversion, there are several questions one would be wise to ask, and also several issues to be borne in mind. The first question is, what is the resolution? In 1-D inversions, this might be represented by horizontal bars on a solution plot, denoting the FWHM of the averaging kernels. Fuller information, at selected points, can be provided by seeing the averaging kernels themselves. In 2-D, it is again feasible only to show kernels at selected targets; or one might display radial width and latitudinal width of kernels as functions of position as plots in their own right. At least in the (one hopes) substantial region where the kernels can be well localized, such measures are an adequate substitute for seeing the kernels themselves. Of course, one also wants to ask, what is the uncertainty ('error bar') on the solution? The formal error at a point, which is a measure of how much the solution at that point could change if the data errors ϵ_i had a different realization, is just $(\sum c_i^2 \sigma_i^2)^{1/2}$ (under the assumption of independent errors) and this can be included as error bars or bands on 1-D plots or can be plotted in its own right. But it should be borne in mind that the assumptions about the size or independence of the data errors may be wrong (see below). Also the error bars are only formal: because the data are integrals, the true Ω could take arbitrarily large (positive or negative) values over a sufficiently small region and not affect the data perceptibly, so that the true solution at a single point has in a sense infinitely great uncertainty (e.g. Genovese et al. 1995). This problem is resolved if one thinks of the solution as a localized average, which is really what the inversion provides. Also, even if the

original data errors are independent, the solution at different points x_0 and x_1 say are different linear combinations but of the same data, so that the errors $\sum c_i \epsilon_i$ at the two locations are correlated (e.g. Howe & Thompson 1996). Howe & Thompson found that, for least-squares and optimally localized averages inversions for rotation, the lengthscale over which errors are positively correlated is similar to the resolution length as defined by the width of the averaging kernel's main peak. Thus it is difficult to know whether a feature in the solution at the limit of resolution is (a) a true feature with that scale; (b) a smaller-scale feature blurred out by the finite resolution of the inversion; or (c) an artifact of the data noise, because the error at neighbouring points in the solution are correlated. (For inversions for structural differences, e.g. for density differences, errors can be correlated over essentially the whole of the solar interior, which introduces further problems of interpretation.) The possibility also exists that features in the solution are due to systematic errors in the data, which we have explicitly ignored in our assumptions about the statistics of the ϵ_i .

6. Sources of uncertainty in current inversions

A quick survey through the literature of recent helioseismic results will reveal that, although the results agree in broad terms, when one looks at fine details there are discrepancies (e.g., Tomczyk et al. 1995; Thompson et al. 1996; Schou et al., 1997b). What has been achieved in making deductions about the solar interior is remarkable, but as the data become more and more precise we want to use that improvement by making more ever more subtle inferences. Thus it is pertinent to ask what is the reason that inversion results are discrepant, in whatever respect. Korzennik et al. (1997) identified three particular sources of uncertainty in current inversions: (a) the inversion results are obtained by different methods and/or different choices of trade-off parameter values; (b) the inversions use data from different instruments and/or observations; and (c) the inversions use data that has been reduced by different peak-bagging pipelines (viz the reduction from observations to finally identifying the modes and quantifying the frequencies and/or frequency splittings). Let us consider each of these briefly.

That applying different inversion methods to the same data yields different looking results should not be a surprise, in the light of the discussion in the preceding section. The data-fitting methods at best yield but one example of a solution that fits the data, so there is no surprise if two such solutions differ, provided the difference between the two solutions lies in the annihilator (see Gough 1985; Gough & Thompson 1991). But even if this is not the case, and of course methods like optimally localized averages don't even purport to find a solution that 'fits' the data, we may interpret the two solutions at a given point as averages of the true solution weighted by the averaging kernels for the respective methods. Because the averaging kernels inevitably differ, the solutions are taking different averages of the true underlying function and if the solutions at a given point differ it does not necessarily mean that the solutions are inconsistent. The solutions are only truly inconsistent if there is no possible underlying Ω which, when convolved with the averaging kernels for the two methods, would give the respective solutions. This is hard to assess, but one possible test when comparing the solution of a data-fitting method with the result of an optimally localized averages approach would perhaps be to see whether the former solution, when convolved with the optimally localized averaging kernels, produced a result that was similar to the optimally localized averages solution. I am not aware that this has ever been done.

Of course the way in which the two methods handle data errors might also produce different results, but if the different averaging kernels are again taken into account, then the difference in the two solutions at a given point should be consistent with the error budgets, $\sum c_i \epsilon_i$ from each of the two methods. Of course, the two errors are correlated and this must be taken into account when making the comparison, but that is easily done if the coefficients c_i are known and the statistical properties of the data errors are as we have assumed. These kinds of intercomparisons have at best only been done to a very limited extent, but Hare and Hounds exercises – where different inversion methods are applied to the same artificial data – are also useful in helping understand the differences between results of different inversion methods, and these have been used in several studies.

The second source of uncertainty is that data come from different instruments or observations. Different instrumental characteristics might lead to systematic differences in the time series. Although a comparison of the time series from the three SOHO experiments shows that they are indubitably looking at the same Sun (Toutain et al. 1997), it is possible that the differences that do exist might lead to systematic errors in the frequency determinations and in the inversions. The actual solar frequencies also change with time. Thus observations from different epochs should only be combined in a single inversion with great caution. In terms of our eq. (1), these sources of uncertainty can be viewed as translating into the errors ϵ_i having a systematic component. (Alternatively, one could view the fact that Ω is time-varying as being an instance of a systematic error in the solution arising from the inadequacy of the underlying model inherent in eq. 1.) One way in which progress can be made on eliminating this source of uncertainty is a careful study of the systematic differences between frequencies and splittings from different instruments or observations.

The third source of uncertainty is that the data is reduced through different peak-bagging pipelines. The peak-bagging is itself a form of inversion, and has similar potential susceptibilities to incorrect assumptions about the errors on the data put into the pipeline. It also in general uses a model of the line shape which is now thought to be deficient in that it does not take into account the asymmetry of the line profile. And the peak-bagging is a nonlinear optimisation problem in which a given algorithm might fall into a local minimum or otherwise be sensitive to an initial guess. Assuming that all the internal checks one can think of have been done, it is highly desirable that systematic comparisons be made of the frequencies and splittings coming out of different peakbagging pipelines, particularly if they can be applied to exactly the same input data. Such comparisons are being undertaken.

I should not end, though, on the problems with inversion! This paper has been a review of some aspects of the inversion methods, and it is a fact that the methods have been developed in helioseismology to an impressive degree. More importantly, the inversions of solar data are revealing the structure and dynamics of the solar interior in ways that a few years ago would have seemed unimaginable. It falls however to other authors in this volume of proceedings to discuss those exciting results.

Acknowledgements

I gratefully acknowledge the co-operation of the SOI Internal Rotation Team, specifically the authors of Schou et al. 1997a and in particular Rachel Howe and Jørgen Christensen-Dalsgaard, for helpful comments and for allowing me to quote the resolution values in Section 4, which are based on their work.

REFERENCES

- Brodsky, M. A. & Vorontsov, S. V., 1988. In *Proc. IAU Symposium No 123, Advances in helio- and asteroseismology*, p. 137 – 140, eds Christensen-Dalsgaard, J. & Frandsen, S., Reidel, Dordrecht.
- Charbonneau, P., Tomczyk, S., Schou, J. & Thompson, M. J., 1997. *Astrophys. J.*, in press.
- Christensen-Dalsgaard, J., Gough, D. O. & Thompson, M. J., 1991. *Astrophys. J.*, **378**, 413 – 437.
- Christensen-Dalsgaard, J., Schou, J. & Thompson, M. J., 1990. *Mon. Not. R. astr. Soc.*, **242**, 353 – 369.
- Christensen-Dalsgaard, J., Duvall, T. L., Gough, D. O., Harvey, J. W. & Rhodes Jr, E. J., 1985. *Nature*, **315**, 378 – 382.
- Genovese, C. R., Stark, P. B. & Thompson, M. J., 1995. *Astrophys. J.*, **443**, 843 – 854.
- Gough, D. O., 1984. *Phil. Trans. R. Soc. London, Ser. A*, **313**, 27 – 38.
- Gough, D. O., 1985. *Solar Phys.*, **100**, 65 – 99.
- Gough, D. O., 1986. In *Seismology of the Sun and the distant Stars*, p. 125 – 140, ed. Gough, D. O., Reidel, Dordrecht.
- Gough, D. O. & Thompson, M. J., 1991. In *Solar interior and atmosphere*, eds Cox, A. N., Livingston, W. C. & Matthews, M., p. 519 – 561, Space Science Series, University of Arizona Press.
- Howe, R. & Thompson, M. J., 1996. *Mon. Not. R. astr. Soc.*, **281**, 1385 – 1392.
- Korzennik, S. G., Thompson, M. J., Toomre, J. & the GONG Internal Rotation Team, 1997. To appear in *Proc. IAU Symposium 181 'Sounding solar and stellar interiors'*, eds Provost, J. & Schmider, F. X. (Kluwer, Dordrecht).
- Kosovichev, A. G., 1996. *Astrophys. J.*, **469**, L61 – L64.
- Pijpers, F. P. & Thompson, M. J., 1992. *Astron. Astrophys.*, **262**, L33 – L36.
- Schou, J., Antia, H. M., Basu, S., et al., 1997a. *Astrophys. J.*, submitted.
- Schou, J. & the SOI Internal Rotation Team, 1997b, these proceedings.
- Shibahashi H., 1988. In *Proc. IAU Symposium No 123, Advances in helio- and asteroseismology*, p. 133 – 136, eds Christensen-Dalsgaard, J. & Frandsen, S., Reidel, Dordrecht.
- Thompson, M. J., Toomre, J., et al., 1996. *Science*, **272**, 1300 – 1305.
- Tomczyk, S., Schou, J. & Thompson, M. J., 1995. *Astrophys. J. Lett.*, **448**, L57 – L60.
- Toutain, T., Appourchaux, T., Baudin, F., et al., 1997. *Solar Phys.*, in press.