THE COOKIE CUTTER TEST FOR TIME-DISTANCE TOMOGRAPHY OF ACTIVE REGIONS

S. G. Korzennik

Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA, 02138

ABSTRACT

Time-Distance helioseismology holds the most promising potential for high resolution three dimensional tomography of active regions. Namely by analyzing the time propagation anomalies resulting from acoustic waves propagating through and interacting with an active region, we should be able to derive enough diagnostic information to map the flows and any sound speed anisotropies below the region surface as such a region develops and evolves, and maybe eventually predict its emergence. Initial inferences from time anomalies computed over and around active regions have shown tantalizing results, but guarded scepticism ought to be held as the formalism of the methodology is still being developed and the nature of the interactions in the magnetized region – especially near the surface – are poorly understood and thus not yet modeled.

I present here the results of a conceptually simple test to validate the putative fast sound speed plume derived below the well-studied active region of June 1998. The idea is to remove from the input set used in the inversion some of the time anomalies computed using observations taken at the surface inside the active regions, measurements that might have been affected by interactions not modeled in the kernels used in the inversion. To validate this test, I first computed simulated data sets based on a similar sound speed perturbation plume and resulting from a direct forward computation. The acoustic ray approximation was used to derive the sensitivity kernels.

I compare how well the underlying sound speed profile is recovered or not in these simulations as observables are removed over circular patches of various sizes centered on the active region (hence the so called "cookie cutter" test). I then present how the inversion behaves when the same test is applied to actual data. Two distinct inversion techniques were used, including one that allows me to compute the corresponding resolution kernels.

The validity of the inferences are discussed in light of the results of these simulations as well as the resulting profile of the corresponding resolution kernels.

Key words: Sun: helioseismology - Sun: Oscillations.

1. INTRODUCTION

Tantalizing results suggesting the presence a fast sound speed "plume" below an isolated sunspot have been derived by Kosovichev *et al.* [1]. This sound speed anisotropy below an active region was inferred from time-distance time propagation anomalies computed using MDI high resolution observations taken on June 18th 1998, while a well isolated circular sunspot crossed the high resolution field of view of the instrument.

While if real such result is quite exciting and would challenge any model of an active region, one ought to exert some guarded scepticism on these initial inferences. The formalism of time-distance inversion is still being developed, especially in active regions. Let us keep in mind that the nature of the interactions in the magnetized region – especially near the surface – is poorly understood and thus not yet modeled.

One is led to wonder whether such inference is real and one is thus motivated to develop a validation test. The validation test I propose here is to remove from the input set used in the inversion some of the time anomalies computed using observations taken at the surface inside the active regions, measurements that might have been affected by interactions not modeled in the kernels used in the inversion. In practice a small set of observables are removed over circular patches of various sizes centered on the active region – hence the so called "cookie cutter" test – prior to carrying out the inversion.

2. THE HYPOTHESIS & METHODOLOGY

The hypothesis I wanted to test is whether the inferred sound speed anisotropy could result from some sort of surface contamination of the measurements of the time anomalies over the active region.

The methodology to test this hypothesis I used is based on the following argument: one should get similar results by dropping the potentially suspicious measured time anomalies.

I carried out inversions using the complete data set and

four separate cutout sets, each with an increasing radius for the cutout area (all centered over the active region).

I also generated simulated data to test the methodology: a forward computation using the same kernels as for the inversion was carried out. I then applied the same methodology on the simulations and the actual data.

3. THE INVERSION METHODS

I used two inversion methods to be sure that the result of such a validation test is not affected by some behavior specific to an inversion method. A robust result will emerge if one sees the same trends when using either method.

The first method is well known, efficient, and used by numerous investigators. The other method is less efficient but allows the computation of the resolution kernels. For both methods and all cases, I used simple sensitivity kernels based on acoustic ray propagation theory.

3.1. Method 1: Iterative Least-Squares (LSQR)

The so-known LSQR method is an iterative conjugate gradient least-squares method [2]. The method is fast and easy to implement, but due to its iterative nature there is no explicit solution, hence one cannot, in practice, derive the corresponding resolution kernels.

Therefore the appropriateness of the choice of the trade-off coefficient for this method – a trade-off between resolution and error magnification – cannot be objectively validated by the corresponding resolution kernels.

3.2. Method 2: Truncated Single-Value Decomposition (TSVD)

The TSVD is a slower method, based as its name suggests on a single-value decomposition (SVD). Fortunately, the basic sensitivity matrix for the time-distance inverse problem is very sparse so implementing in this context an SVD remains manageable. To increase efficiency the method can be, and has been, parallelized.

The main advantage of the TSVD method is that there is an explicit solution, hence one can compute the corresponding resolution kernels and use them to validate objectively the choice of the trade-off coefficient.

3.2.1. Details of the TSVD method

The least-squares normal equations of any inverse problem are:

$$A^T\, y \;=\; A^T A\; x \tag{1}$$

$$\begin{aligned}[m \times n][n] &= [m \times n][n \times m][m] \\ [m] &= [m \times m][m]\end{aligned}$$

that can be rewritten as:

$$b = Mx \tag{2}$$

where y represents the observables (the time anomalies), x the underlying model (the sound speed anisotropy), and A the sensitivity matrix. Hence M is a square matrix ($M = A^T A$), whose *formal* solution is given by:

$$x = M^{-1} b \tag{3}$$

but in practice M is near-singular and has no numerically stable inverse.

Any matrix can be factorized using a single-value decomposition (SVD), from which the formal inverse can be derived. Namely M can be rewritten as:

$$M = U \Lambda V^T \tag{4}$$

and its formal inverse as:

$$M^{-1} = V\,(1/\Lambda)\,U^T \tag{5}$$

since $U^T U = 1$ and $V^T V = 1$.

The matrix Λ is a diagonal matrix (with diagonals elements λ_i), and thus its inverse $(1/\Lambda)$ is trivial (the diagonal matrix with elements $1/\lambda_i$), but is finite only if all the single values are non-zero.

In most numerical implementations the matrix M will be near-singular, resulting in some very small single values. These will naturally contribute to a numerically unstable inverse $(1/\Lambda)$.

In the truncated single-value decomposition method, an estimate of the solution is computed by truncating the matrix $(1/\Lambda)$, setting $1/\lambda_i \rightarrow 0$ for $\lambda_i < \epsilon\, \lambda_{\max}$.

The choice of ϵ controls the amount of regularization, hence it acts as the trade-off between resolution and error magnification.

An estimate of the solution for a value of ϵ is thus given by:

$$x_\epsilon = M_\epsilon^{-1} b = M_\epsilon^{-1} A^T y \tag{6}$$

Standard error propagation gives us the uncertainty on the inverse solution as:

$$\sigma_x^2 = |M_\epsilon^{-1} A^T|^2\, \sigma_y^2 = |U\,(1/\Lambda)\,U^T A^T|^2\, \sigma_y^2 \tag{7}$$

Of course, when the uncertainty of observables, σ_y, are known, the problem is rescaled by dividing both the observables y and the columns of the matrix A by the respective uncertainties, and thus adequately weighting the reliable information present in the observables.

The explicit formulation of the solution (Eq. 6) combined with Eq. 1 lead to a straightforward expression for the resolution kernels, R, since:

$$x_\epsilon = M_\epsilon^{-1} A^T y = M_\epsilon^{-1}\, A^T\, A\, x \stackrel{\rm def}{=} R\, x \qquad (8)$$

where x_ϵ is the computed estimate of the solution while x represents the *real* underlying solution.

The resolution kernel, R, is thus "simply":

$$R(\vec{x}_{\rm Target}, \vec{x}) = M^{-1}_{\epsilon,{\rm Target}}\, A^T\, A \qquad (9)$$

The resolution kernel indicates the effective spatial resolution and actual localization of the computed solution. It is the *key* element to judge objectively the estimate of the solution resulting from the choice of ϵ. It not only indicates the effective spatial resolution of the solution, but whether the estimated solution is actually localized at the target location and whether some of the "features" present in the solution are real or due to non-local contaminations.

Note also that in practice, since $(1/\Lambda)$ ends up being truncated, I compute only $m_{\rm eff} < m$ single values and vectors. Also, since A is sparse — while $M = A^T A$ is not — I use a sparse SVD solver and store A^T as a sparse matrix.

In the absence of observable uncertainties the problem is translation invariant and I can reduce the size of the SVD and compute it only over an appropriate sub-space in (x, y) of the volume analyzed.

The current implementation has also been parallelized, using a parallelized sparse SVD algorithm to make use of multiple CPUs execution speed up combined with an evenly balanced storage distribution, as large problems lead to very large memory requirements.

4. INPUT DATA SET

The data & simulations span $17.4^{\rm o}$ or 209.5 Mm; the cutout radii are:

	R_1	R_2	R_3	R_4
radius [Mm]	6.96	13.92	20.88	27.84
fraction kept [%]	99.68	98.68	96.94	94.53

5. RESULTS FROM SIMULATIONS

I have computed solutions for two simulations for a set of trade-off values and using both inversion methodologies. The first simulation is a simple faster than its surrounding plume (case #1) while the second one consists of a contrasted plume, with a slower than it surrounding volume just below the surface and a faster than its surrounding volume deeper in. The simulated topologies are shown in the lower left panel of all figures showing inversion results for the simulations.

Figure 3 shows inversion results for case #1 using the LSQR method, while results for the same case but resulting from the TSVD inversion method are shown in Fig. 4. Inversions for case #2 are shown in Fig. 5 and Fig. 6 for the LSQR and TSVD methods respectively.

In all four figures one does see clear indication of the presence of a localized sound anisotropy in *each* panel. Namely for this range of trade-off parameters, for all four cases of cutouts and for both methods, one recovers some indication of the presence of a sound speed plume in the simulation. The actual shape of the inferred plume and the present or not of artifacts are modulated – as one should expect – by the value of the trade-off coefficient and the extent of the cutout.

6. RESULTS FROM ACTUAL DATA

The two inversion methods where applied to the actual data for the same trade-off values and the same cutoff radii as the simulations shown in the previous section. The resulting inferences are shown in Figs. 7 and 8. In these figures the lower left panels show inferred values for the corresponding trade-off parameter without any cutout.

We do recover the plume seen in [1], but notice how in Figs. 7 and 8 the "plume-like" feature is seen only in the lower left panels: for the cases that include all the data or only small cutouts and for larger values of ϵ (*i.e.*, cases that include fewer eigenvalues and thus have a lower resolution).

7. RESOLUTION KERNELS

The TSVD method allows me to compute the corresponding resolution kernels. Figure 9 shows the resulting resolution kernels at four depths for three trade-off parameter values that correspond to the TSVD inversions shown in the top, middle and bottom panels and for the case where all the observables were included (*i.e.*, no cut).

Figure 10 shows the resulting resolution kernels for only one depth for the same cases (trade-off parameter values and cutout extend) as the figures showing TSVD inversion results (Figs. 4, 6, and 8).

8. CONCLUSIONS

The cookie cutter validation test strongly suggests that the sound speed "plume" below the active region results

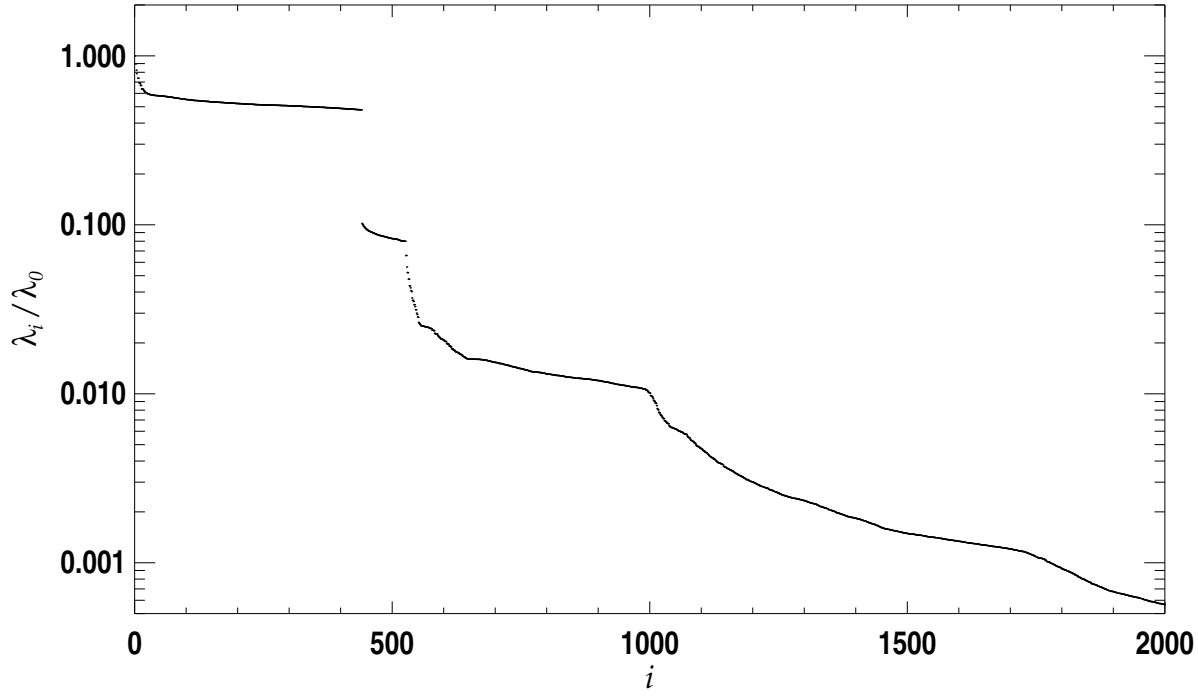

Figure 1. The 2,000 largest eigenvalues – out of 9,800 – for the single-value decomposition of the sound speed perturbation problem corresponding to travel time anomalies computed for the active region of June 18, 1998 using MDI high-resolution mode. Note the succession of large jumps of the eigenvalues, indicative of a good truncation location.

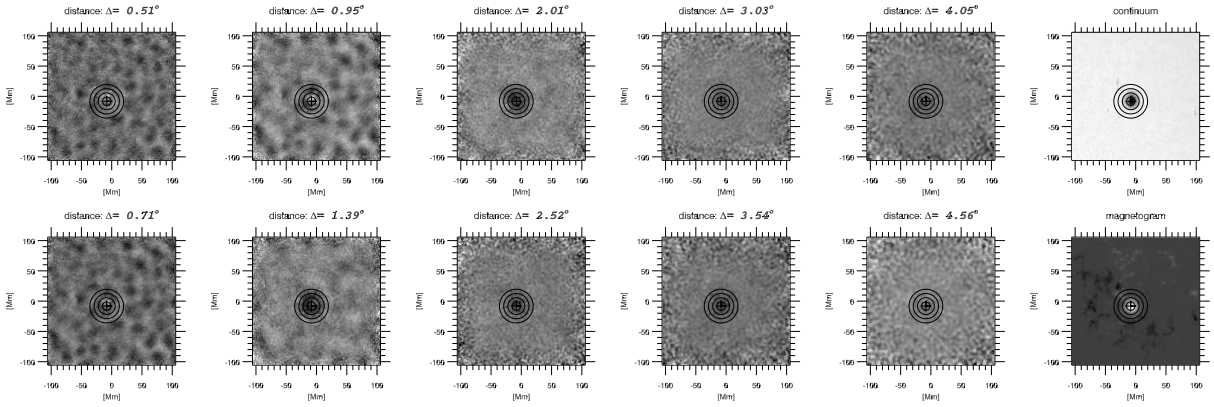

Figure 2. Input data set for 10 distances shown versus position. The two leftmost panels show the corresponding active region extent as seen in the corresponding continuum image and magnetogram. The four circles, in each panel, illustrate the extent of the four cylindrical cuts.

from a surface contamination. Indeed, we see a consistent pattern in the inferred solutions when using two distinct inversion methods: the plume does not survive much of a cut while the simulations clearly show that some indication of the plume should survive the cookie cutter test. This result is consistent with the conclusions from Rajaguru *et al.* [3].

The resolution kernels give us us the actual localization (width and depth) of the solution. They indicate what trade-off values correspond to acceptable resolutions, and provide a handle on the effective depth. The cutouts widen the kernel – as one would expect – but the kernel location remains comparable, especially the effective depth.

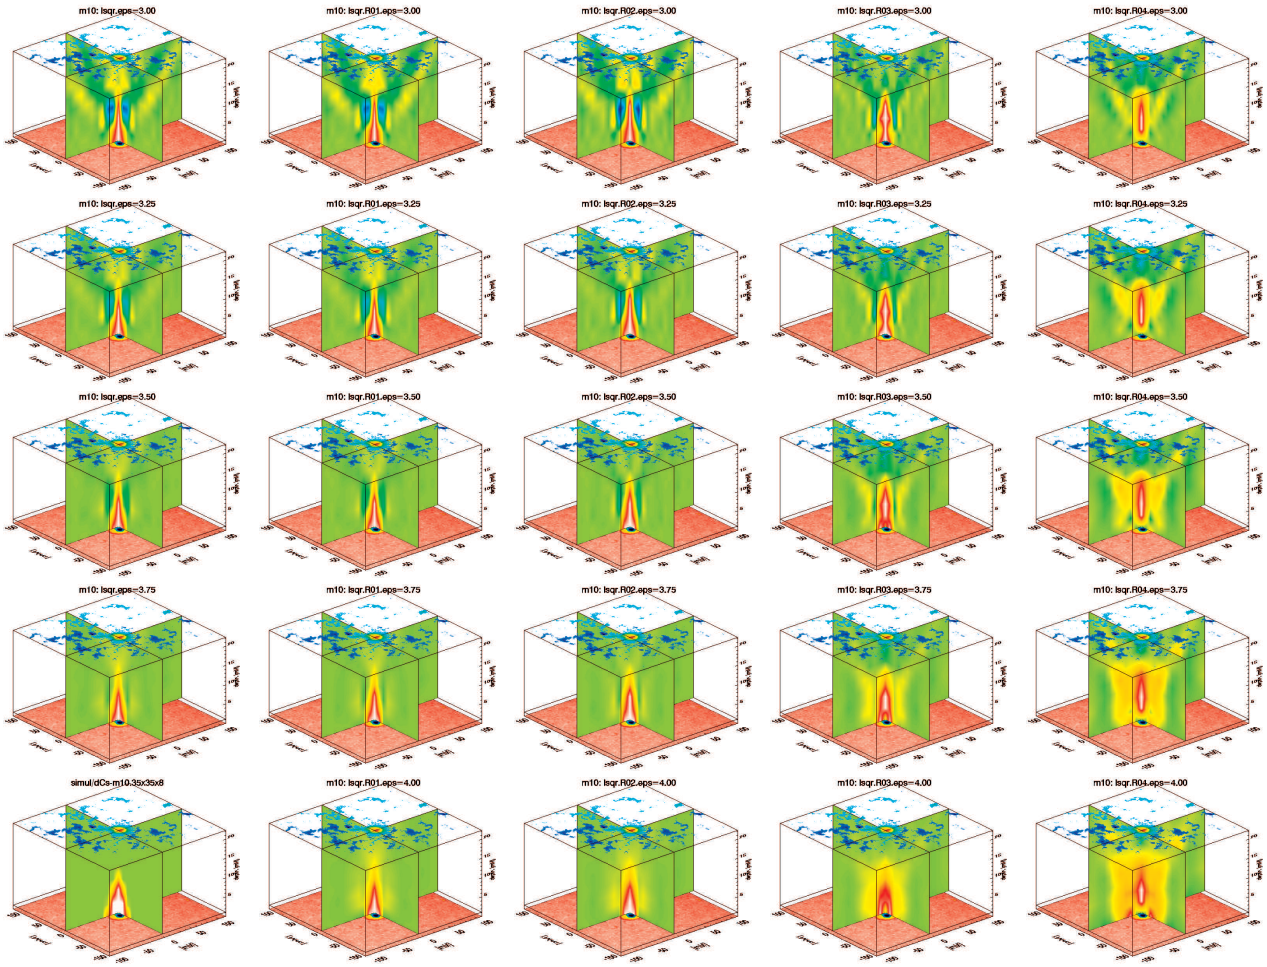

Figure 3. Inferences for one simulation, using the iterative conjugate gradient method (LSQR) inversion methodology. The five columns correspond to including all the data (leftmost column), and to data with cuts at the four radii of increasing value (radius increases from left to right). The five rows correspond to different values of the trade-off parameter – trade-off between resolution and error magnification – with the solution resolution increasing from bottom to top. The lower left panel is the input topology for that simulation, or the true solution. In these views, the solar surface is at the bottom of the data cube.

ACKNOWLEDGMENTS

The Solar Oscillations Investigation (SOI) involving MDI is supported by NASA grant NNG05GH14G at Stanford University. SOHO is a mission of international cooperation between ESA and NASA. This work was supported by NSF grant ATM–0318390 and NASA grant NAG5–13501.

REFERENCES

[1] Kosovichev, A. G., Duvall, T. L., & Scherrer, P. H. 2000, Sol. Phys., 192, 159

[2] Paige, C. C., Saunders, M. A., 1982, ACM Trans. Math. Software, 8, 43

[3] Rajaguru, S. P., Birch, A. C., Duvall, T. L., Jr., Thompson, M. J., & Zhao, J. 2006, Ap. J., 646, 543

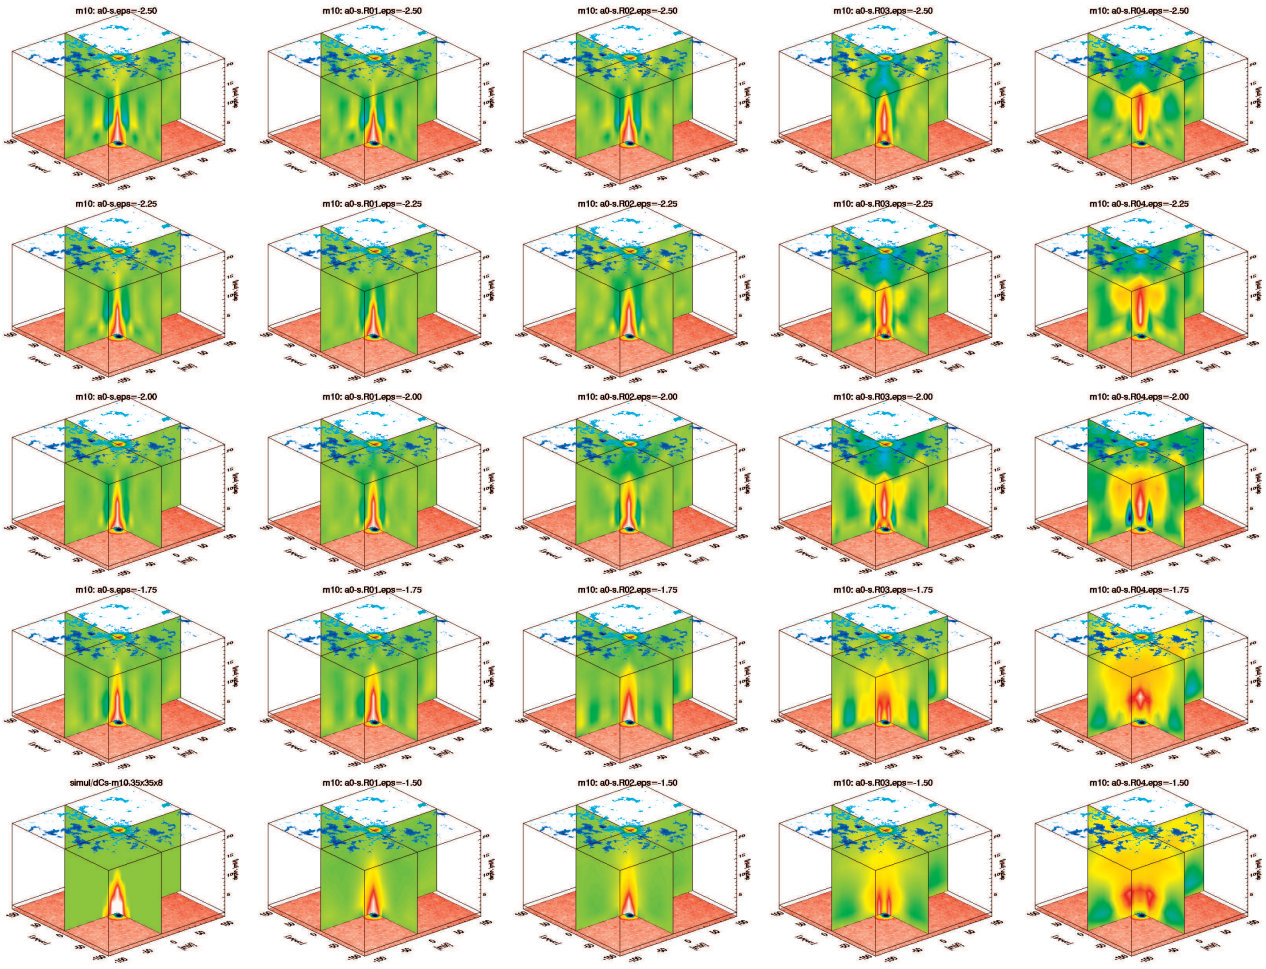

Figure 4. Inferences for the same simulation as in Fig. 3, but using the truncated single-value decomposition (TSVD). The five columns correspond to including all the data (leftmost column), and to data with cuts at the four radii of increasing value (radius increases from left to right). The five rows correspond to different values of the trade-off parameter – trade-off between resolution and error magnification – with the solution resolution increasing from bottom to top. The lower left panel is the input topology for that simulation, or the true solution. In these views, the solar surface is at the bottom of the data cube.

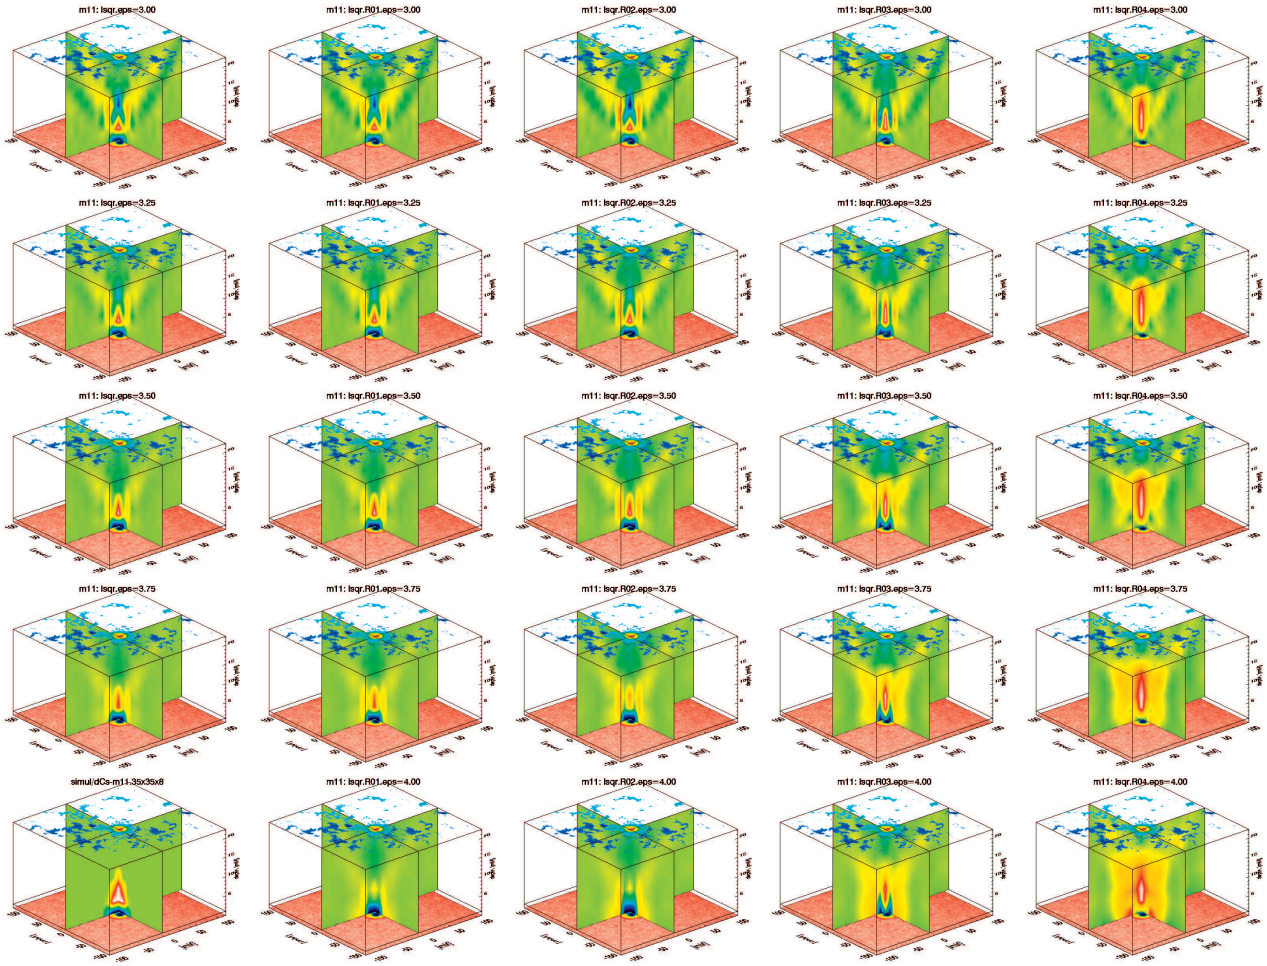

Figure 5. Similar to Fig. 3, but for a slightly different simulation (see lower left panel: contrast near the surface.) Columns correspond to including all data or 4 different cuts; rows correspond to different values of the trade-off parameter – with the solution resolution increasing from bottom to top.

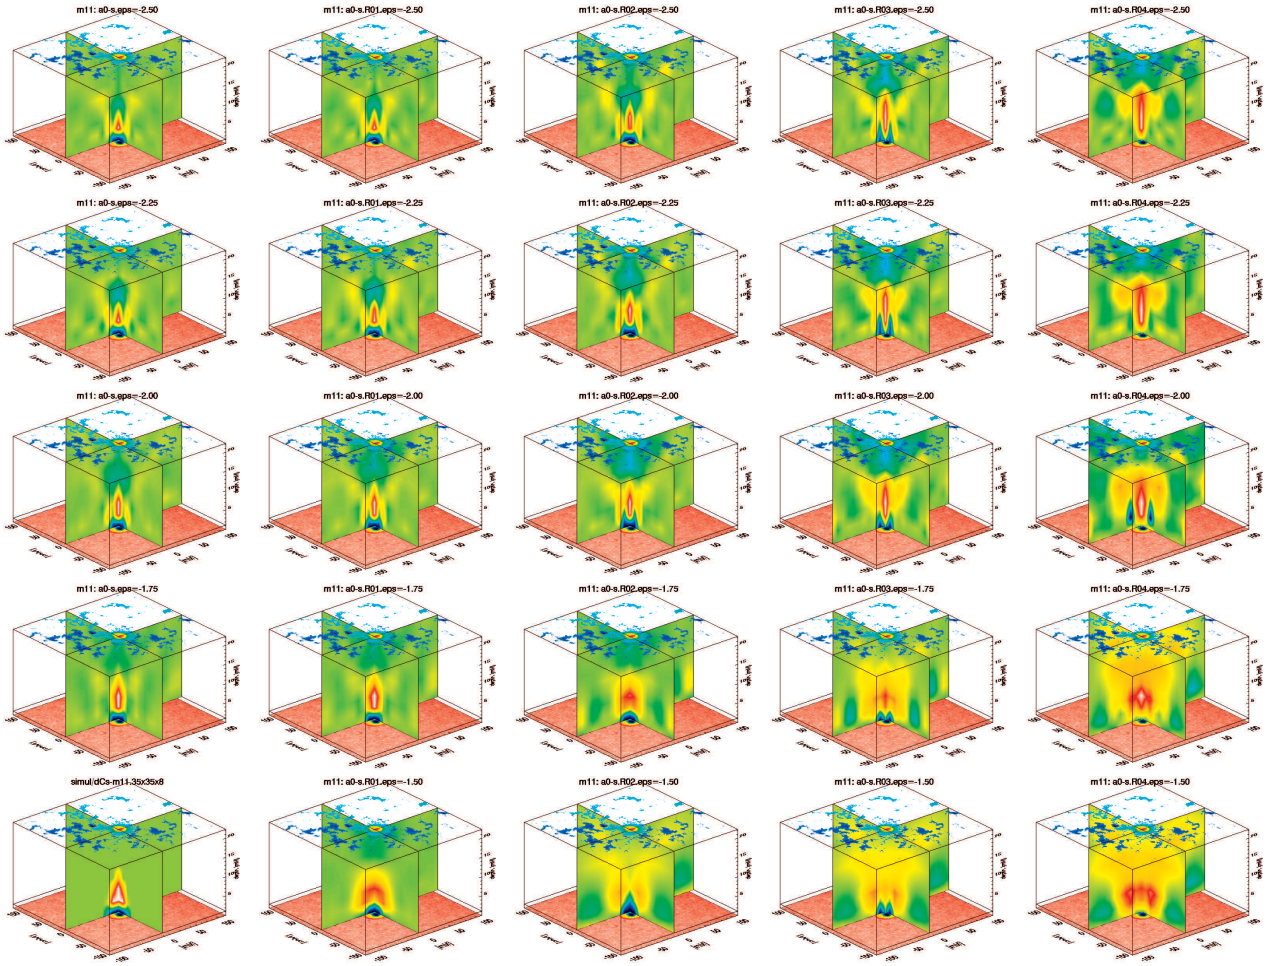

Figure 6. Similar to Fig. 4, but for a slightly different simulation (see lower left panel: contrast near the surface.) Columns correspond to including all data or 4 different cuts; rows correspond to different values of the trade-off parameter– with the solution resolution increasing from bottom to top.

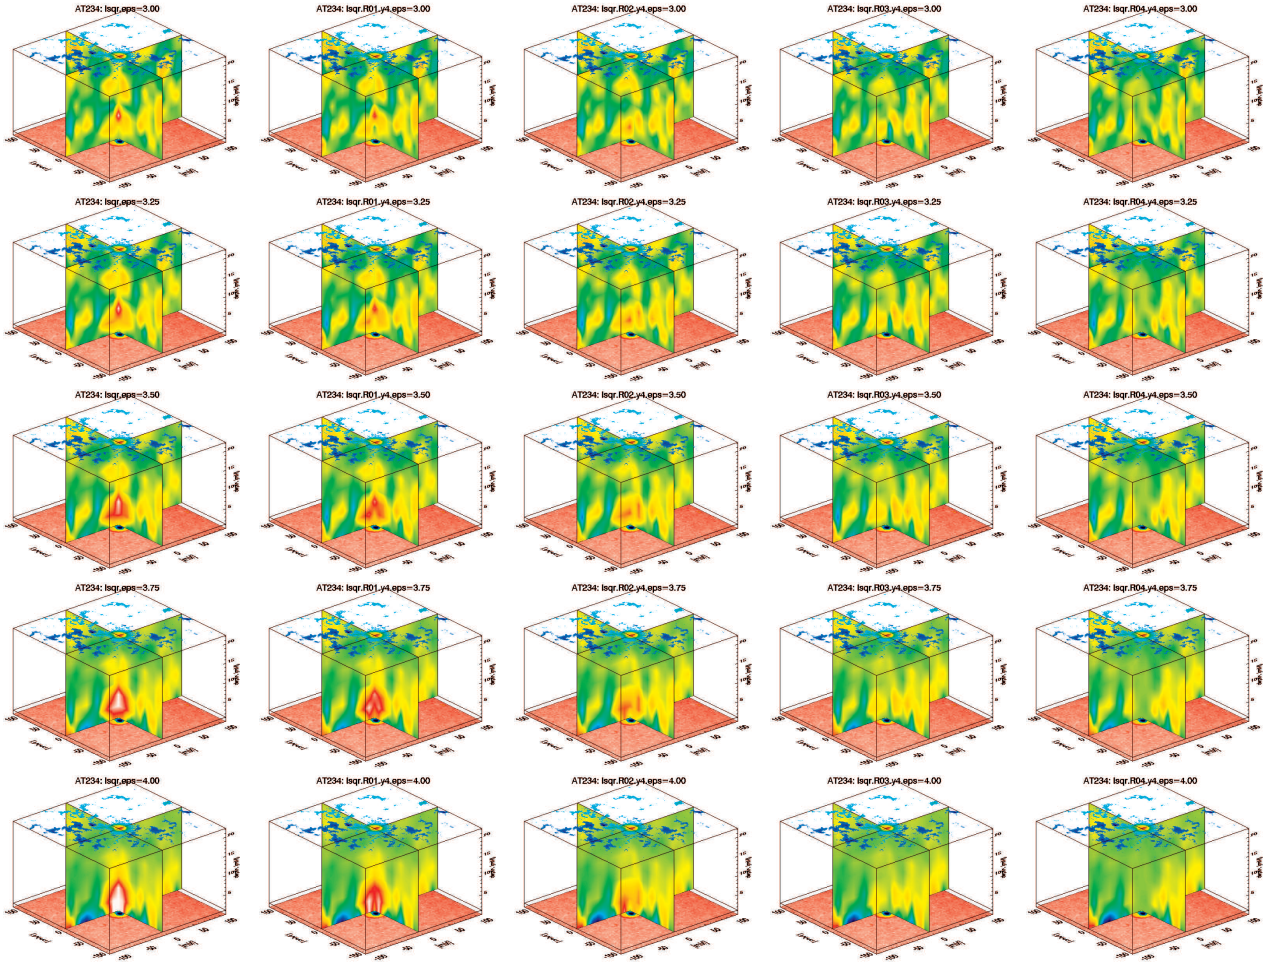

Figure 7. Inferences for the same methodology as Figs. 3 & 5 (LSQR), but using actual data (inferred values for all panels since there is no known true solution.)

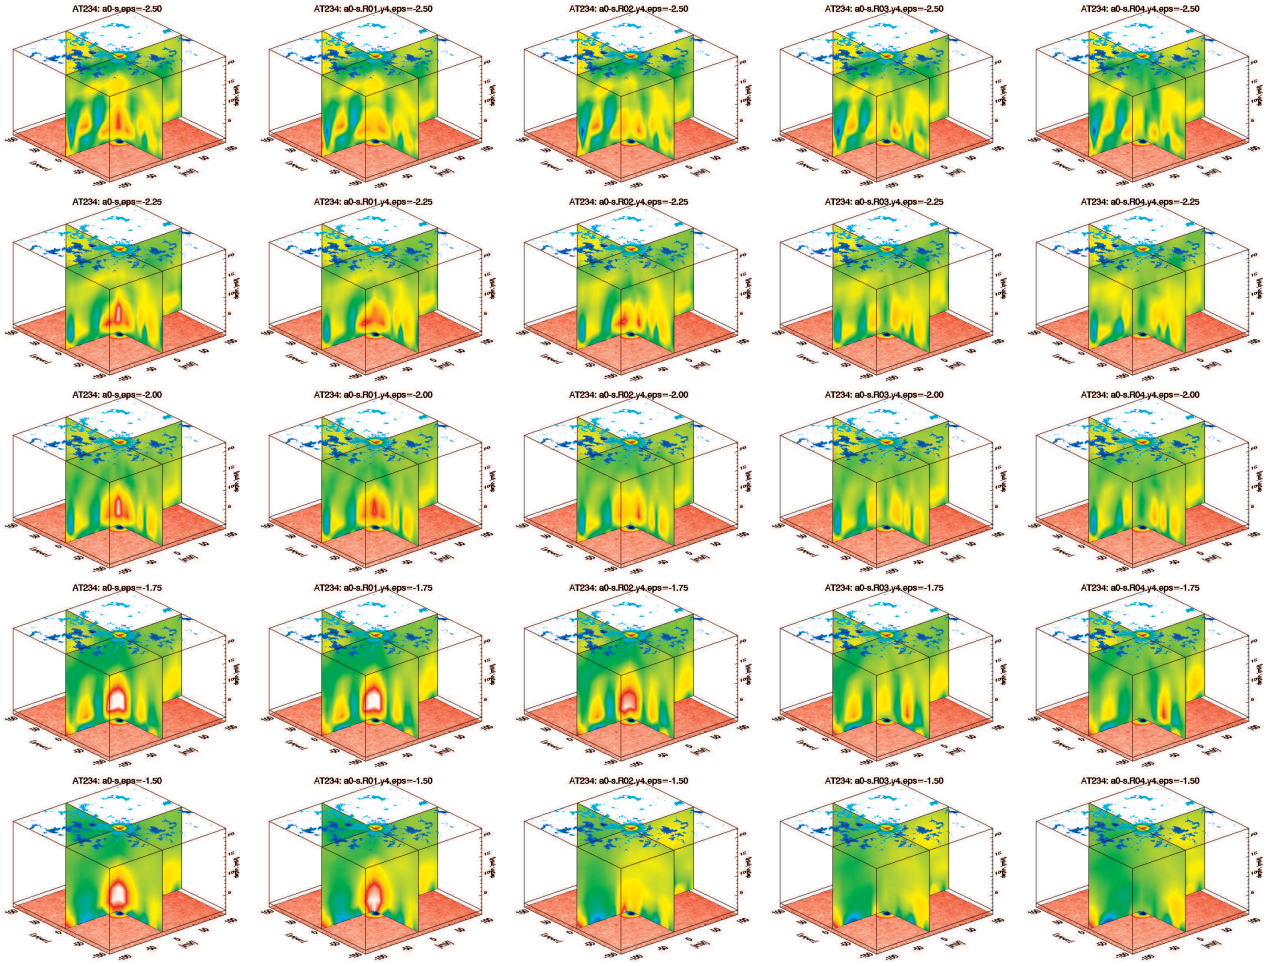

Figure 8. Inferences for the same methodology as Figs. 4 & 6 (TSVD), but using actual data (inferred values for all panels since there is no known true solution.)

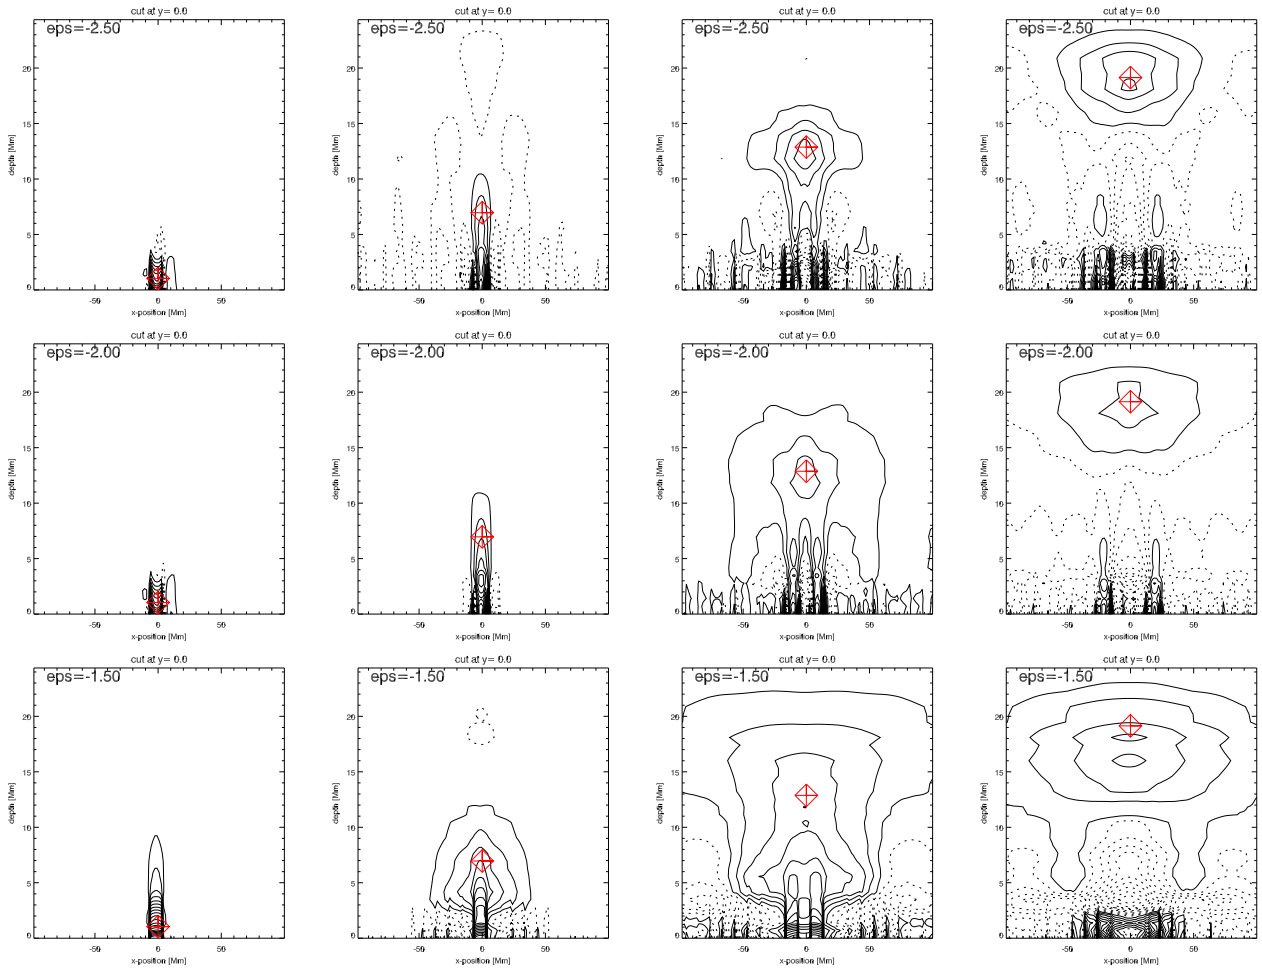

*Figure 9. Resolution kernels, cut at $y = 0$ shown as a function of position and depth, for five depths (columns) and three trade-off values (rows). The kernels shown here correspond to a case that uses all the observables (*i.e.*, no cut). The red symbol indicates the target location.*

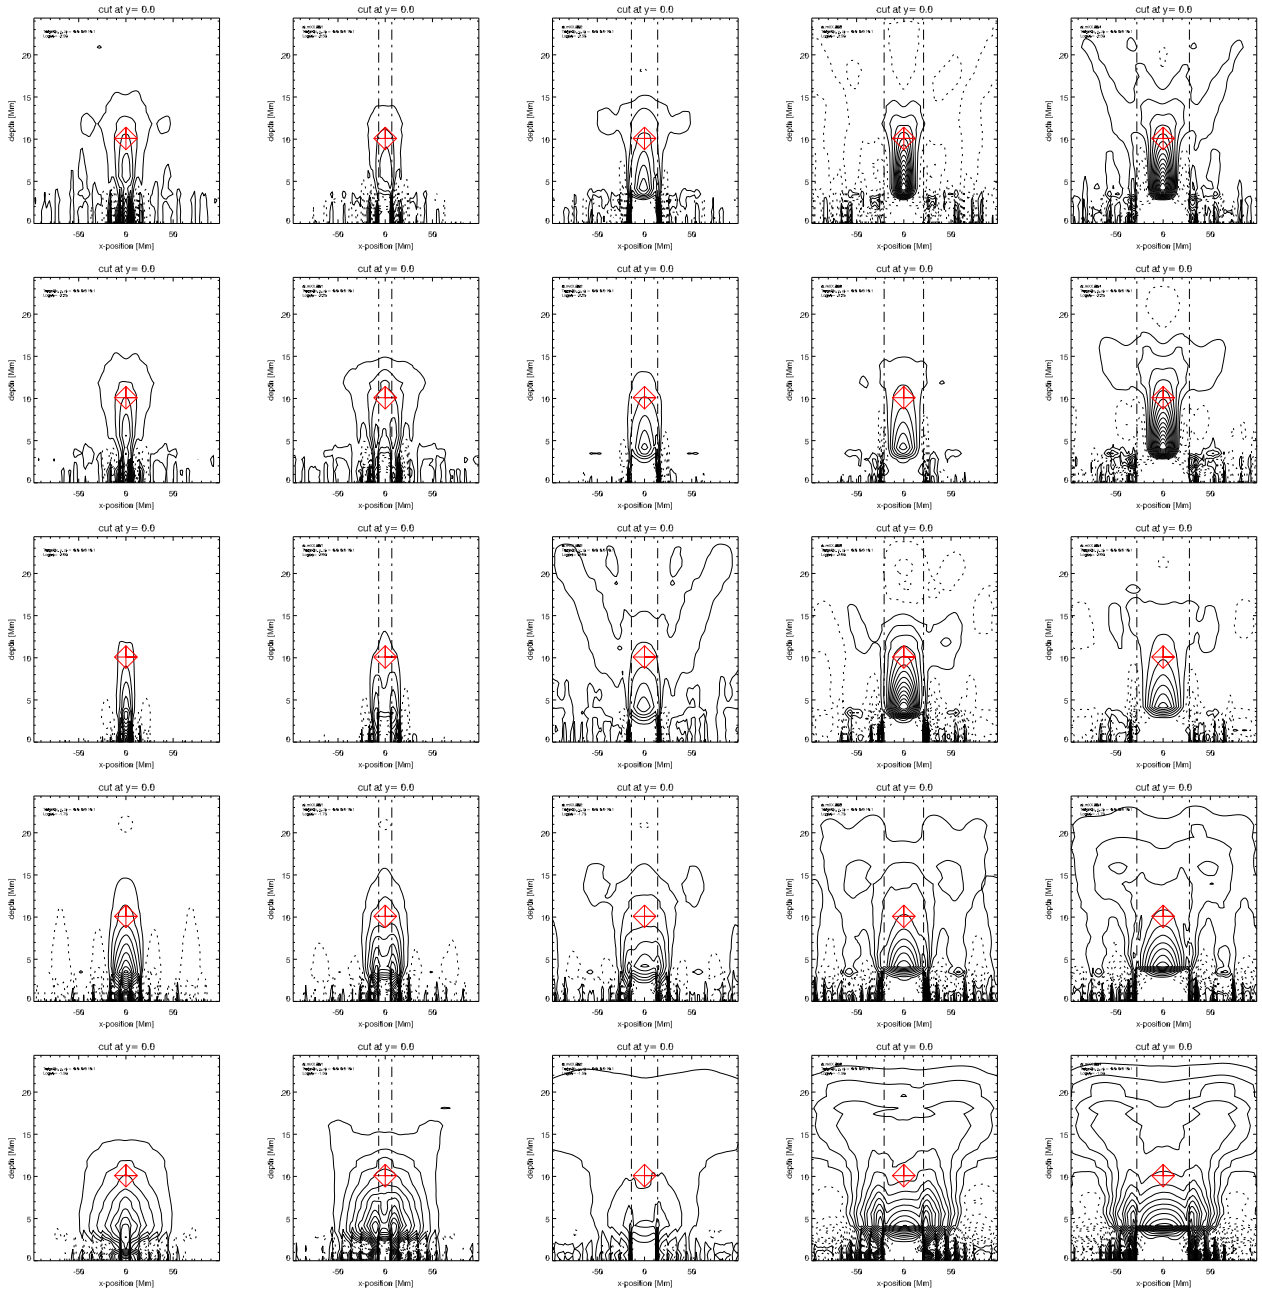

Figure 10. Resolution kernels, cut at $y = 0$ shown as a function of position and depth, for one given target depth only (indicated by the red symbol). The panels show the resolution kernel for five input set cases (columns: all, and 4 cuts) and five trade-off values (rows), corresponding to the inferences shown in Figs. 4, 6 and 8. The vertical lines illustrate the radial extend of the cutouts.