$\mathcal{AIPS}$ Memo 118 Revised

Modeling Spectral Cubes in $\mathcal{AIPS}$

Eric W. Greisen

October 1, 2025

Abstract

$\mathcal{AIPS}$ has done Gaussian fitting along the x-axis of image cubes with task XGAUS since the 1980s. That task was overhauled in 2013 to be much easier to use and much more capable. In like fashion, tasks ZEMAN and RMFIT were developed. The former fits the standard leakage and scaling terms for Stokes V cubes, including a new option to do this for each of the Gaussians found by XGAUS. The latter fits polarization models to Stokes Q and U cubes, using the output of Faraday Rotation Measure Synthesis ($\mathcal{AIPS}$ task FARS) to assist with initial guesses. The models can contain multiple components each with a polarization flux, angle, rotation measure, and rotation measure "thickness." The present memo will describe the functions of these tasks in some detail with numerous graphical examples. A revision of this memo discusses changes made for the 31DEC15 release and two new tasks to plot spectra with model fits and a number of tasks which make visibility and image model files. The November 2017 revision includes new display options in RMFIT. The April and July 2025 revisions add display and edit options and correct display and pixrange setting. The September 2025 revision contains changes made to allow more than 8 components in XGAUS and ZEMAN. [1]

1 Introduction

XGAUS was written, and re-written, a long time ago in the hopes that converting a spectral cube, with many spectral channels at each spatial location, into one or more components each with a set of images of peak, center, width, and integral would simplify the data presentation. One could even hope that the separate Gaussians fit might have some separate physical reality. The task could fit a single Gaussian to each spatial position without much human guidance, but detailed human interactivity was required for more complex spectra. Use of the task involved an endurance contest between the human and the power company, requiring both luck and skill on the part of the human. The task had no way to back up and fix fits generated at previous spatial positions and either wrote out all of the fit parameters in images or discarded all work done on the cube. Clearly, this was not a desirable situation and would-be users tended to prefer fitting routines in other packages such as GYPSY.

There has always been some interest in fitting rotation-measures to polarization images taken at different frequencies. Until the advent of the Jansky VLA however, the images available for fitting rotation measures were at a modest number of scattered frequencies and could generally be fit by rather simple algorithms. This is no longer the case and better fitting algorithms were required. Initially there was considerable interest in a method known as Faraday Rotation Measure Synthesis.[2] The method was implemented in $\mathcal{AIPS}$ by Leonid Kogan as task FARS with supplementary tasks and additions by the author. However, it was found that the "synthesized beam width" in this method was too wide to separate most scientifically interesting cases, even if the data were taken to frequencies as low as 1 GHz. Data from lower frequencies would be needed to produce better component discrimination, but most radio sources are completely depolarized at these low frequencies. A more direct fitting method is therefore required.

[1] Color is used to indicate changes made to this document: for the April, July and September, 2025 revisions.
[2] Brentjens, M. A. and de Bruyn, A. G. 2005, *Astron. & Ap.*, 1217.

The author undertook a re-write of XGAUS in 2013 in order to make it a much more useful and user-friendly task and to provide a software framework for a new task to do polarization fitting. The suggestion was then made to create a task to fit standard Zeeman-splitting parameters as well. The following Memo attempts to describe in detail the function of all three of the tasks that have resulted. This memo has been revised for the 31DEC15 version of these tasks and does not attempt to describe differences with the 31DEC13 and 31DEC14 versions, in which these tasks first appeared. The 2025 modifications have added useful new capabilities. Two tasks, AGAUS and ZAMAN, have appeared in the 31DEC17 release of *AIPS* to fit Gaussians and Zeeman splitting in absorption spectra. They are described in *AIPS* Memo 122, which was also revised in 2025.

Two tasks have been written to prepare plot files from the fits done by these tasks. In 31DEC25, there is a new task to copy the extension files of these tasks in order to change the number of components modeled. There are also tasks in *AIPS* which modify uv-data files to add model components both spatially and spectrally. In addition there are tasks which add polarization and spectral models to images. These will be described in a final section.

2 Gaussian fitting: XGAUS

Gaussian fitting is the process of determining the peak brightness, center position, and width of a number of one-dimensional Gaussians, plus a spectral baseline, which best matches a one-dimensional slice of input data. XGAUS will do this for every voxel in a three-dimensional image cube, taking each row as the one-dimensional slice. It then produces images over the second and third input axes of the fit parameters. Normal usage has the first axis of the cube as a spectral axis, either in frequency or velocity units, and the second and third axes are celestial coordinates. This usage will be assumed in the later discussion, but it is not required. Mathematically, the function fit is

$$T(x,y,z) = A + Bx + \sum_{1}^{\mathrm{NGAUSS}} T_i e^{-\alpha(x-x_i)^2/\sigma_i^2} \tag{1}$$

where x is the coordinate along the first axis of the image cube, $A(y,z) + B(y,z)x$ is the spectral baseline, $T_i(y,z)$ are the peak values of each of the NGAUSS components, $x_i(y,z)$ are the center channels of the components, $\sigma_i(y,z)$ are the full widths at half maximum of the components and $\alpha = 4\ln(2)$ to impart this meaning to the σ_i.

2.1 Inputs

The usual INNAME *et al.* adverbs define the cube to be fit and the OUTNAME *et al.* adverbs define the (eventual) output name. The INVERS adverb controls which XG table is used by the task, with ≤ 0 meaning a new table. BLC and TRC define the pixel ranges to be used in the current execution, where BLC(1) and TRC(1) control the spectral channels that will be fit and BLC(2), TRC(2), BLC(3), and TRC(3) control the area in celestial coordinates over which the fitting is done. New XG table files are created for the entire input image cube, but are filled with peak brightnesses limited by BLC(1) through TRC(1). Therefore, it would be wise to use as much of the first axis as is reliable when creating a new table. The second and third values of these adverbs then limit the area over which fitting is done during this execution of XGAUS. Adverbs YINC and ZINC control the stride taken in the first pass through the cube; a second pass will then fit all voxels not fit in the first pass. FLUX controls which spectra will be fit during this execution; all positions with 3 consecutive channels averaging above FLUX will be fit. The initial guess for the linear baseline is now always zero. Adverb ORDER controls the order of the baseline with < 0 meaning none, 0 meaning a constant, and ≥ 1 meaning a constant plus a slope. DOOUTPUT controls what files are written — this may be changed interactively (except for one-dimensional images) so leave it zero at this point. Set DOTV> 0 to use TV menus to prompt you. Even when fitting only 1 Gaussian component, it is best to watch what is happening so you should never set this adverb false. DORESID controls whether residuals are plotted on the fit spectra; such plots often provide clues when more Gaussians are needed for the best fits. In general it is best to set PIXRANGE to zero

to see the full range of image values. In 31DEC25, you may change PIXRANGE interactively, but only when the task reaches the edit stage. Cutting off the highest values in order to see weaker features better may make it harder to set initial guesses however since you can't point at the peak of a clipped component. Set LTYPE to your favorite type of labeling, set PIXVAL to zero to see all positions being fit, and leave NITER zero since 100 is more than enough iterations. Set NGAUSS to the number of Gaussians to be fit in this execution. If a new XG table is created, NGAUSS sets the maximum number of components that may appear in that table. In 31DEC25, NGAUSS may be as large as 32, while previous versions limited NGAUSS to 8 and always set the maximum in the table to eight. NGAUSS may be changed if you re-start on a pre-existing XG table but only with less than or equal that original number of Gaussians. A task to get around this limitation has also appeared in 31DEC25. XG2XG will be discussed below. RMSLIMIT is an upper limit for the rms of a fit before the fit is viewed as "failed" which causes the TV and interaction to be turned back on after you have turned it off. You should get a good idea of an appropriate value from your initial uses of XGAUS or from your knowledge of the noise in your data cube.

The task begins by creating an XG table and populating each row with the largest average brightness over three consecutive channels in the corresponding row in the range BLC(1) through TRC(1). Then it reads the table every YINC rows and ZINC planes and, for those with an average brightness greater than FLUX, attempts a fit. Your interaction with this fit will be described below. After the first pass, the task loops over every row and plane fitting those positions which have enough brightness and which have not already been fit. Finally, after all pixels above FLUX have been fit, the task goes into an "edit" mode. It constructs images of each fit parameter and the integral flux of the Gaussian and of the uncertainties in these parameters. In 31DEC25, it also constructs an image of the rms over the x axis for every (y, z) pixel. This image has a value at every pixel unlike the parameter images which have non-blanked values only at the pixels where the fit was regarded as good. You may view these images, select positions explicitly or by their parameter values or rms and revisit the fits of those positions. This stage will be described in detail below.

At any time you may exit the task and then re-start it using the same XG table. Good reasons for doing this include fitting smaller regions with each pass using the appropriate number of Gaussians for that region. Doing small regions which will have similar parameter values helps a great deal with the initial guessing done by the task (mostly using the previous solution). You might also fit the cube initially with a high value of FLUX and then re-start with a lower value to extend the areas fit.

2.2 Fitting

The fitting process starts with a plot of the spectrum with the data and axis labels in graphics channel one (usually yellow) and the initial guess as X's in graphics channel two (usually green). Then you are offered a menu of options, as shown in the present figures, on the TV (DOTV>0). The first spectrum to be fit (from a model image cube for clarity) is illustrated in the top of Figure 1 showing that the first guess for more than one Gaussian is not useful. The menu that appears at this point is

DO FIT	Attempt a fit
RE-GUESS	Interactively enter a new initial guess
BAD	Mark pixel bad and go on
QUIT	Exit XGAUS now

You select a menu option by moving the cursor to the desired option with the mouse and registering that move with the TV by clicking the left mouse button. The selected menu item will change color as shown in the figure. If you press TV "button" D at this point (actually keyboard character D), helpful information about the selected item will appear in your terminal window. If you press one of TV "buttons" A, B, or C (actually keyboard characters A, B, C), the selected function will be performed. The option to QUIT causes the task to quit at this point. You may re-start later. The option BAD will mark this position as failed and go on to the next position. The option DO FIT will cause the task to attempt the non-linear fit with the current initial guess. The selected option in Figure 1 is RE-GUESS which causes the task to change the color of the

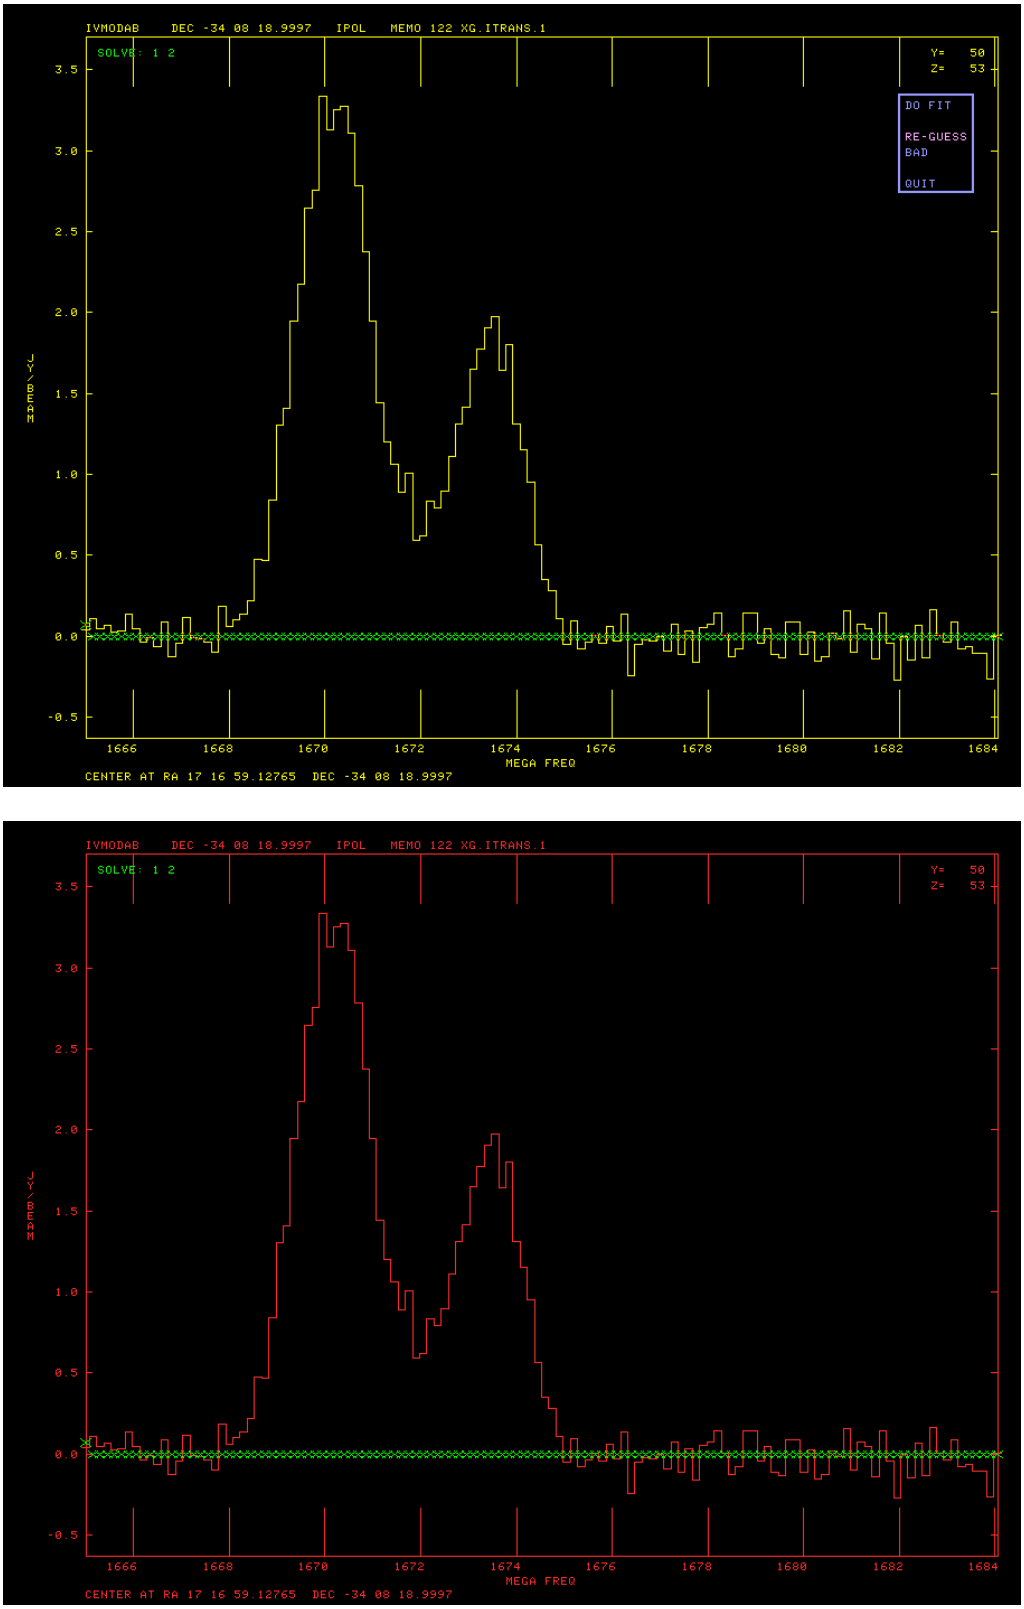

Figure 1: Top: First model spectrum to fit, initial "guess" is all zero. Bottom: better initial guess is requested. Note change of color.

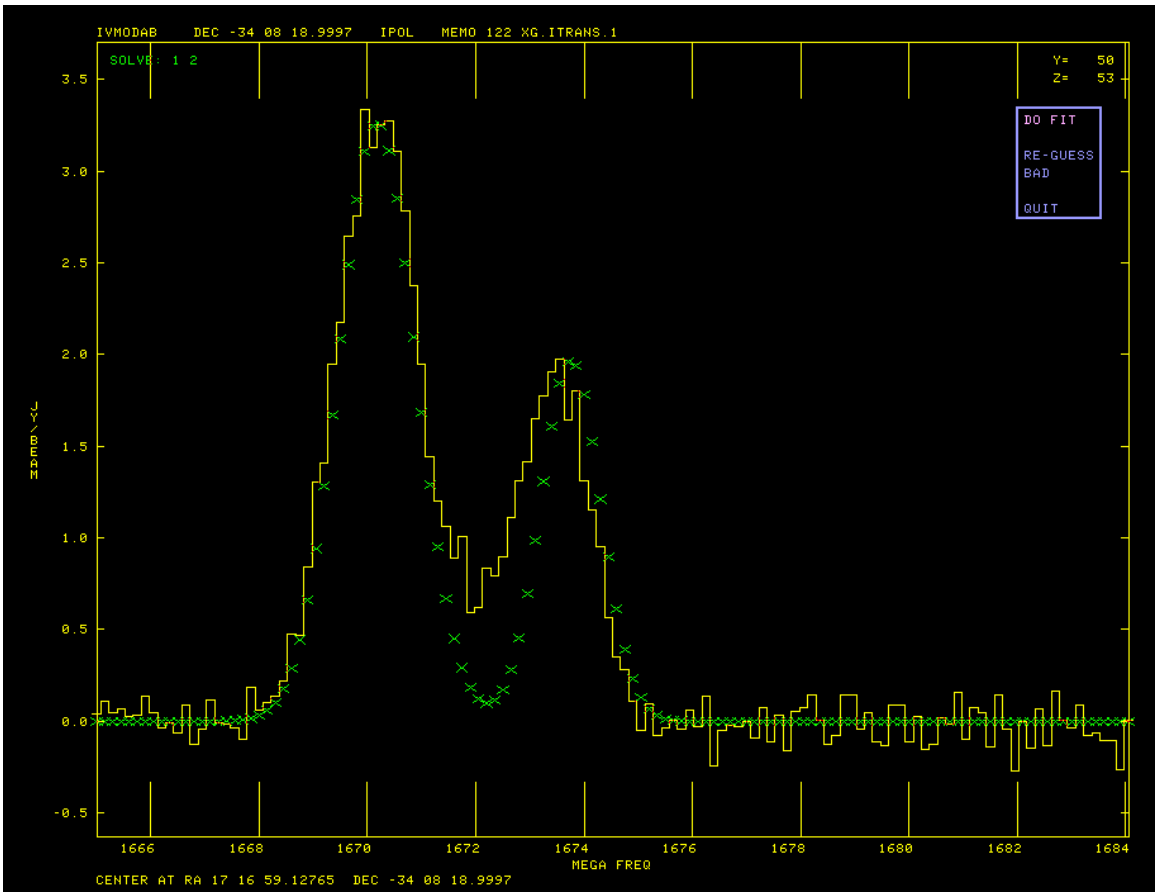

Figure 2: First model spectrum to fit: better initial "guess" has been entered.

plot as illustrated in the bottom of Figure 1. `XGAUS` will then prompt you first to *"Position cursor at center & height of Gaussian component 1"*. Move the cursor to the peak of component 1 and press any TV button. This selects the peak value and center of component 1. Then the task prompts you to *"Position cursor at half-width of Gaussian component 1"*. Move the cursor horizontally to the approximate position of the half-power point of component 1 and press any TV button. The horizontal position of the cursor then sets the initial guess of the full width of the component. These prompts are then repeated for components two through `NGAUSS`. If you do not want to fit a particular component at this position, move the cursor outside the rectangular border line (*i.e.*, outside the data area of the plot, but still inside the TV) before pressing the TV button for that component; it will be omitted.

After the improved initial guess has been entered, the plot is changed to show the new guess, as illustrated in Figure 2. This guess is good so the `DO FIT` option is selected. After TV "buttons" `A`, `B`, or `C` are pressed the task attempts the fit with the current initial guess. Then the plot is changed with the addition of the fit function in graphics plane 4 (usually cyan) and the residual in graphics plane 7 (usually a dirty yellow). This is illustrated in Figure 3. A different menu appears at this point containing

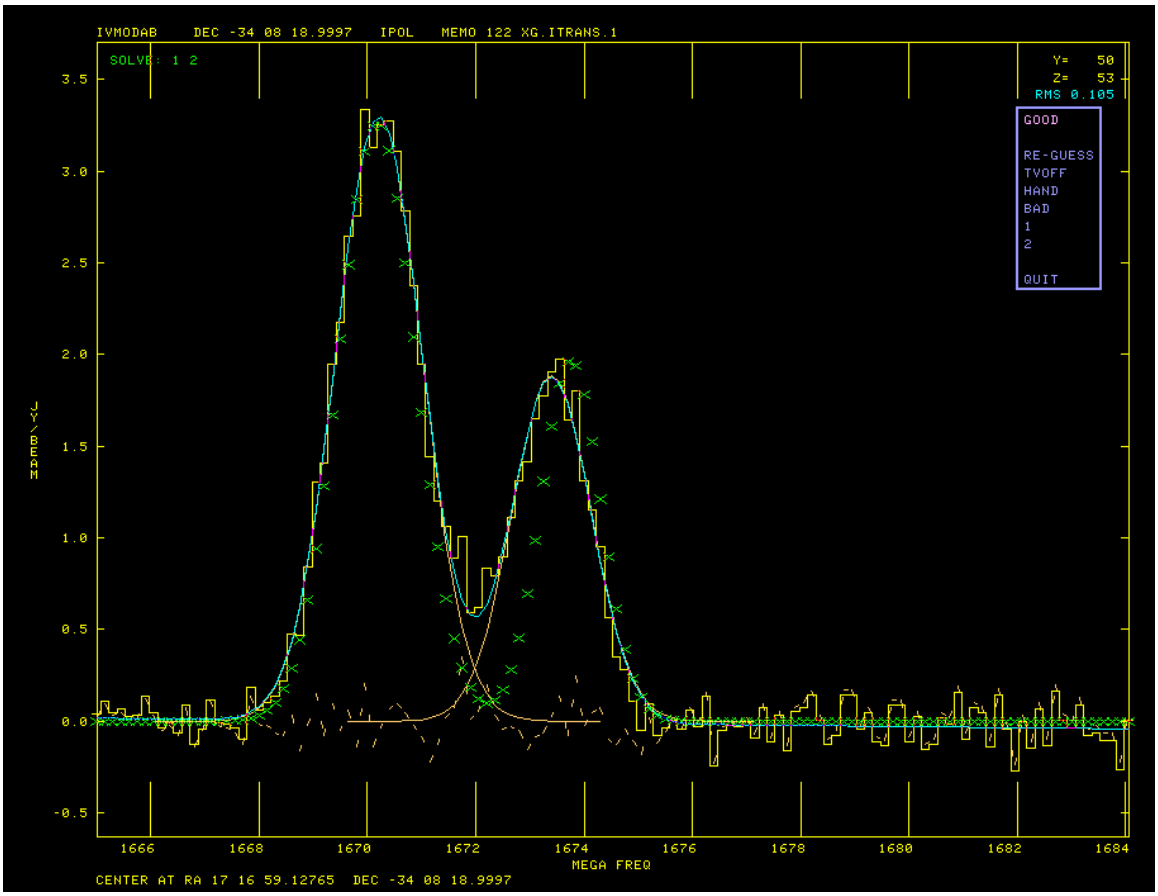

Figure 3: First model spectrum to fit: good fit obtained.

GOOD	Mark solution good and go on
DO FIT	Attempt a new fit
RE-GUESS	Interactively enter a new initial guess
TVOFF	Mark solution good and turn off interactivity
HAND	Enter a solution from the terminal
BAD	Mark pixel bad and go on
1	Interactively enter a 1-Gaussian initial guess
2	Interactively enter a 2-Gaussian initial guess
3	Interactively enter a 3-Gaussian initial guess
4	Interactively enter a 4-Gaussian initial guess
ENTER NG	Interactively enter a number of Gaussians and an initial guess
QUIT	Exit XGAUS now

The option to QUIT causes the task to quit at this point. You may re-start later. The option BAD will mark this position as failed and go on to the next position. The option RE-GUESS will loop back to prompt you for a new guess and repeat the fit. Options 1, 2, 3, 4, ENTER NG will loop back to plot an initial guess with the selected number of Gaussians. The ENTER NG option appears only when the current NGAUSS is more than 4 and prompts you on the terminal to enter a number. (NGAUSS=2 in the current example.) Option HAND will prompt you to enter using the terminal the Gaussian parameters for each component. Enter on one line for each component, the peak value of the Gaussian (in image units), the center (in pixels with respect to the reference pixel), and width (in pixels). Appropriate ranges of parameters in these units can be seen from the display of the current fit values. You may also enter flags to cause one or more parameter values to be

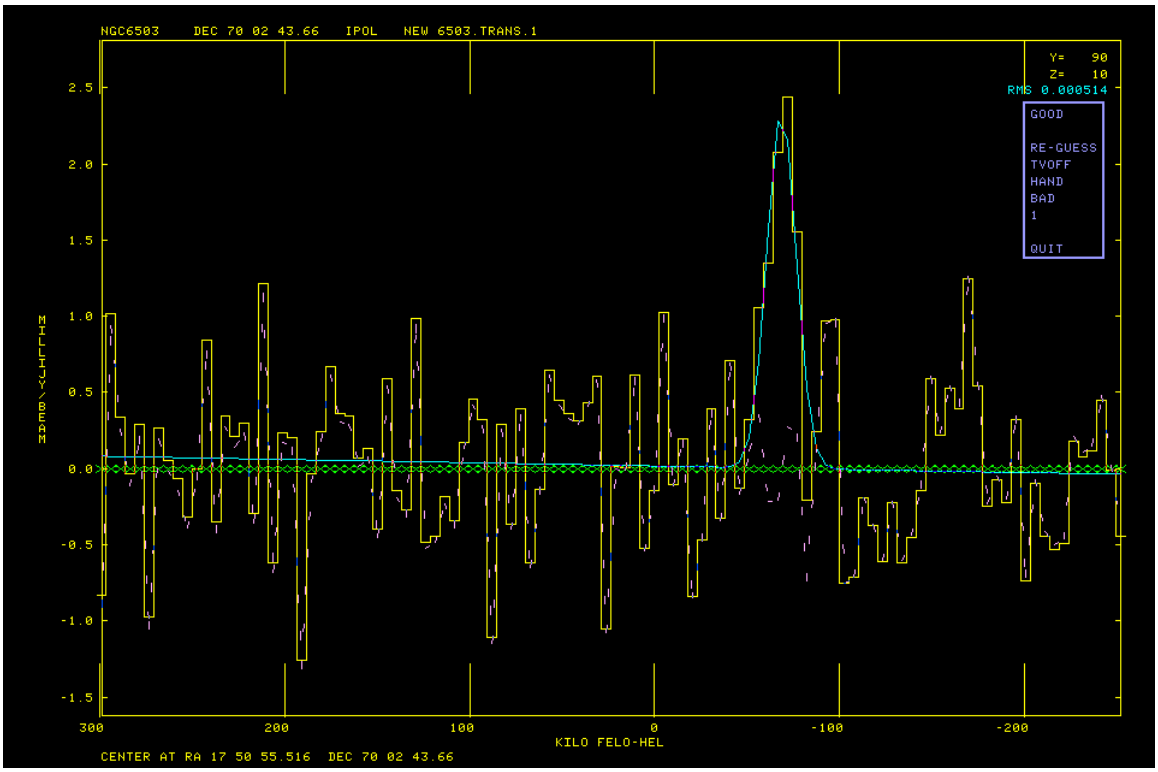

Figure 4: First 1-Gaussian HI spectrum of NGC6503 to fit: good fit obtained automatically.

fixed should you re-fit the current spectrum. The flags are entered after the 3 parameter values, flags ≤ 0 mean to hold the corresponding parameter fixed and omitted flags are taken as 1. `XGAUS` will then repeat the display in Figure 3 to see if you made a good guess. Immediately after a `HAND` operation only, the option `DO FIT` is offered to go back with the hand-entered values as the initial guess for a new fit.

Option `TVOFF` allows you to turn off interactivity, allowing the task to run using its own initial guesses until it finds a completely unreasonable solution or one with an rms greater than `RMSLIMIT`. When that happens, you are shown the offending fit parameters and the task resumes interactivity. Frequently this will drop you immediately into entering a new guess as in the bottom of Figure 1 where the color change should make this apparent. Point to the Gaussian peak and half power following the instructions on the terminal. Sometimes it simply shows something like the top part of Figure 1 and you must select `RE-GUESS` to try to fix things. Option `GOOD` tells the task that you are (reasonably) happy and that it should go on to the next position.

When fitting only a single Gaussian, `XGAUS` makes an initial guess based mostly on finding a real peak in the spectrum. This is quite reliable, so turning off the TV interaction may save a great deal of effort, although there will possibly be bad positions to be fixed up in the next stage of this task. An example of an HI cube, on its first position, is shown in Figure 4.

2.3 Editing and output

Eventually all positions selected by `FLUX` and `BLC` and `TRC` will have been fit. At this point, the task computes images of the fit parameters plus the "flux" (area under the Gaussian) and their uncertainties. An image of the rms over the first axis (after any fit components are subtracted) is also computed in the September 31DEC25 version. This "edit" stage is omitted if the input image is one dimensional, but, in that case, channel flagging is available in the initial fitting of the spectrum. `XGAUS` then offers a lengthy menu of options which

will allow you to view these images and revisit positions that seem to have produced incorrect fits. If NGAUSS > 1, options to swap portions of image n with corresponding portions of image m are also offered. This "edit" menu is illustrated in Figure 5. Note that the size of the signal portion of the model image is quite small. When an appropriate sub-image is selected for display, as was done for this figure, XGAUS replicates pixels in both directions to make the image large enough to see.

There are three kinds of editing implemented here. In the first, the user establishes the parameter extrema which should be viewed as acceptable. The extrema currently set are shown in the title lines. When NGAUSS is greater than nine, only the 9 components surrounding the selected component are displayed. Then, XGAUS may be told to flag all solutions not meeting these criteria, or, more profitably perhaps, to revisit those positions to see if a better fit can be obtained. The other editing methods are similar, but act on a list of pixel positions. These may be entered by typing in values, by clicking on suspect pixel positions in the CURVALUE function described below, or by the BLOTCH LIST function also described below. The contents of the list may be viewed, the solutions at the positions may be flagged, or they may be revisited to attempt for a better solution, or the solutions at the listed positions may be swapped between components n and m. The menu will offer only appropriate swaps between components, thus 1 and 2 in our NGAUSS=2 example. Swapping may be required if XGAUS gets confused as to which component you want to call number 1 and which number 2. Such confusion appears in Figure 5. After the flagging, revisiting, or swapping, the list is cleared. The first column of the menu shows the following:

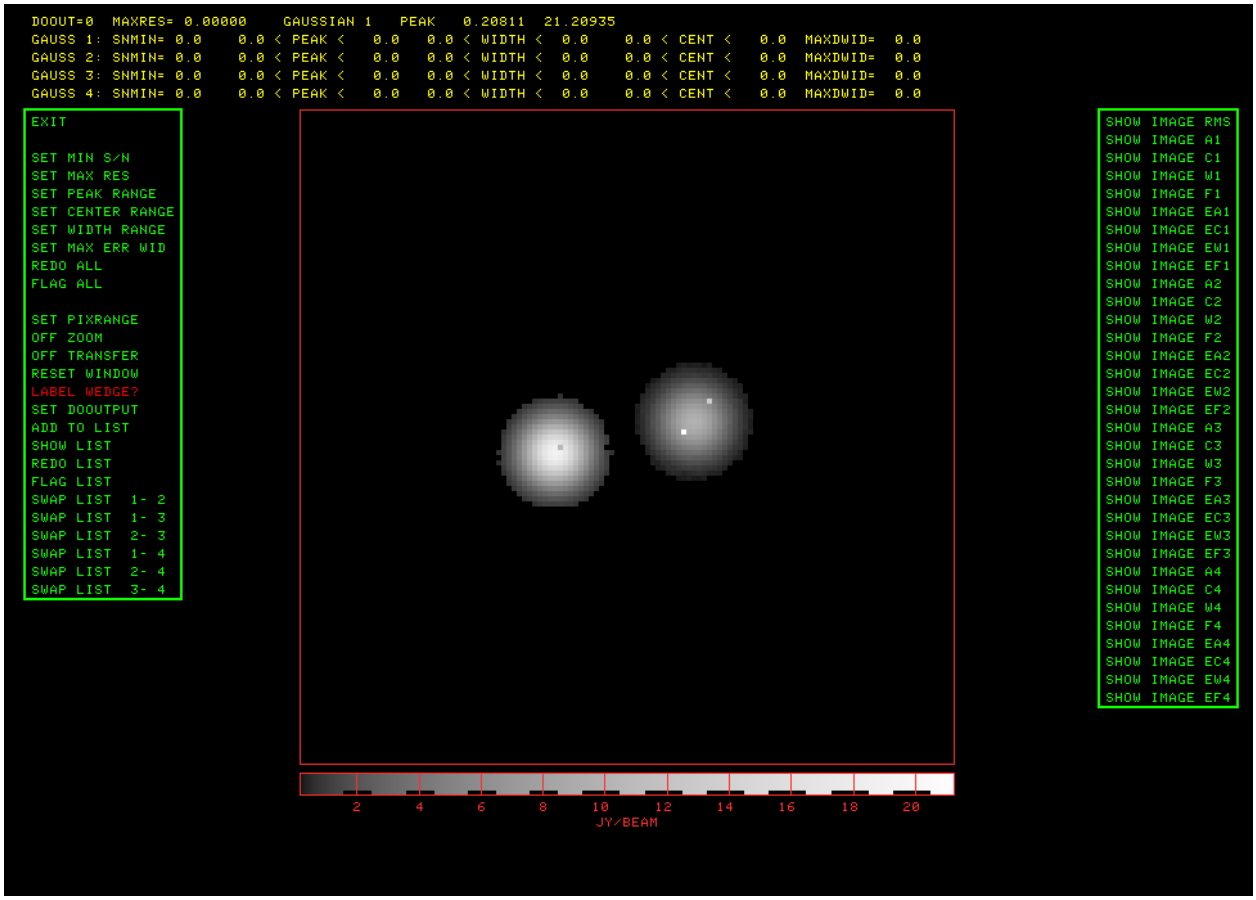

Figure 5: Model source, edit view with `NGAUSS = 4`. Some pixels need swapping or re-doing.

`EXIT`	Exit `XGAUS`, writing output images if `DOOUTPUT` is now > 0.
`SET MIN S/N`	Set minimum amplitude S/N(s) for okay solutions
`SET MAX RES`	Set maximum residual for okay solutions
`SET PEAK RANGE`	Set peak value range(s) for okay solutions
`SET CENTER RANGE`	Set X center range(s) for okay solutions
`SET WIDTH RANGE`	Set width range(s) for okay solutions
`SET MAX ERR WID`	Set maximum error(s) in width for okay solutions
`REDO ALL`	Re-do all solutions which are not okay following the above criteria
`FLAG ALL`	Mark bad all solutions which are not okay
`SET PIXRANGE`	Enter new spectra plot range on terminal
`OFF ZOOM`	Turn of TV zoom
`OFF TRANSFER`	Turn off black & white and color TV enhancements
`RESET WINDOW`	Display full view of current image
`LABEL WEDGE?`	Turn on/off labeling of step wedge
`SET DOOUTPUT`	Increment `DOOUTPUT` in loop 0-3 — with 1 and 3 causing residual images and 2 and 3 causing parameter images to be written on `EXIT`
`ADD TO LIST`	Type in pixel coordinates to add to list
`SHOW LIST`	Display coordinates in list
`REDO LIST`	Re-do solutions for all pixels in list
`FLAG LIST`	Flag solutions for all pixels in list
`SWAP LIST 1-2`	Swap solutions for components 1 and 2 for all pixels in list

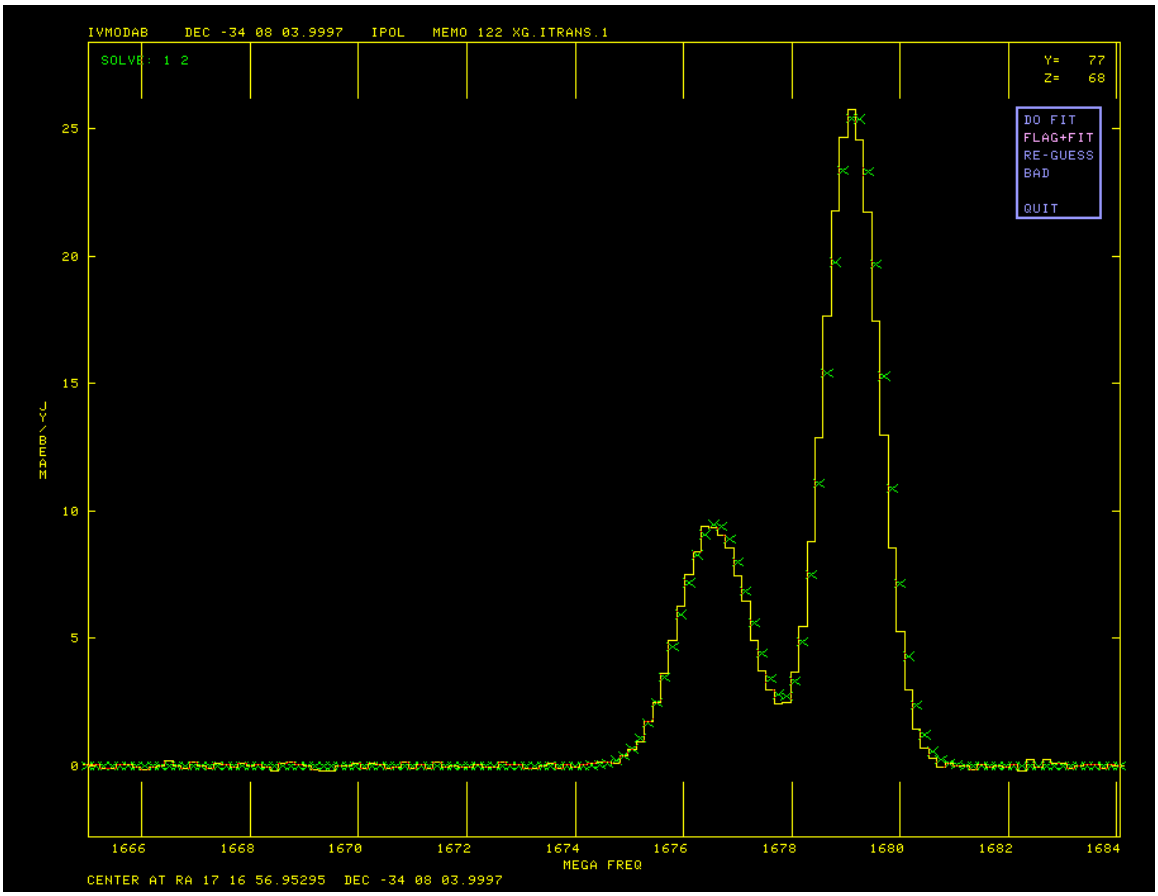

Figure 6: Spectrum to be re-done after doing a RE-GUESS. FLAG+FIT selected for illustrative purposes.

In 31DEC25, XGAUS offers an additional capability in the edit stage.[3] When doing REDO ALL and REDO LIST, the menu shown in Figure 1 and Figure 2 has an additional option FLAG+FIT as shown in Figure 6. When this option is selected a new menu is shown as in Figure 7:

UNDO ALL	Undo all channel flagging
FLAG CHANS	Interactively flag a range of channels
UNDO CHANS	Interactively undo flags in a range of channels
RE-GUESS	Interactively enter a new guess for the Gaussians
DO THE FIT	Leave the flag routine and do the fit

Flagged channels are plotted in a different color than unflagged channels (usually reddish and yellow, respectively). UNDO ALL simply removes all flags in the current spectrum. FLAG CHANS and UNDO CHANS bring up two vertical lines (usually pink in color). The cursor position controls the position of the lines, button A switches between the two, button B sets a channel range and continues, button C sets a range and exits, and button D exits without setting a new range. RE-GUESS takes you to the screen to set a new guess (Figure 1) and then returns to this flag menu. Finally DO THE FIT attempts the Gaussian fit and returns to the menu to accept the result. That screen is illustrated in the bottom of Figure 7. Note the FLAG+REDO option added to this menu in the edit phase. That option takes you back through the sequence just illustrated.

When the XG table is created with fewer than six Gaussians, the second column of the main edit menu contains the rms image plus NGAUSS sets of functions

[3] This option is available at the beginning when the input image is one-dimensional.

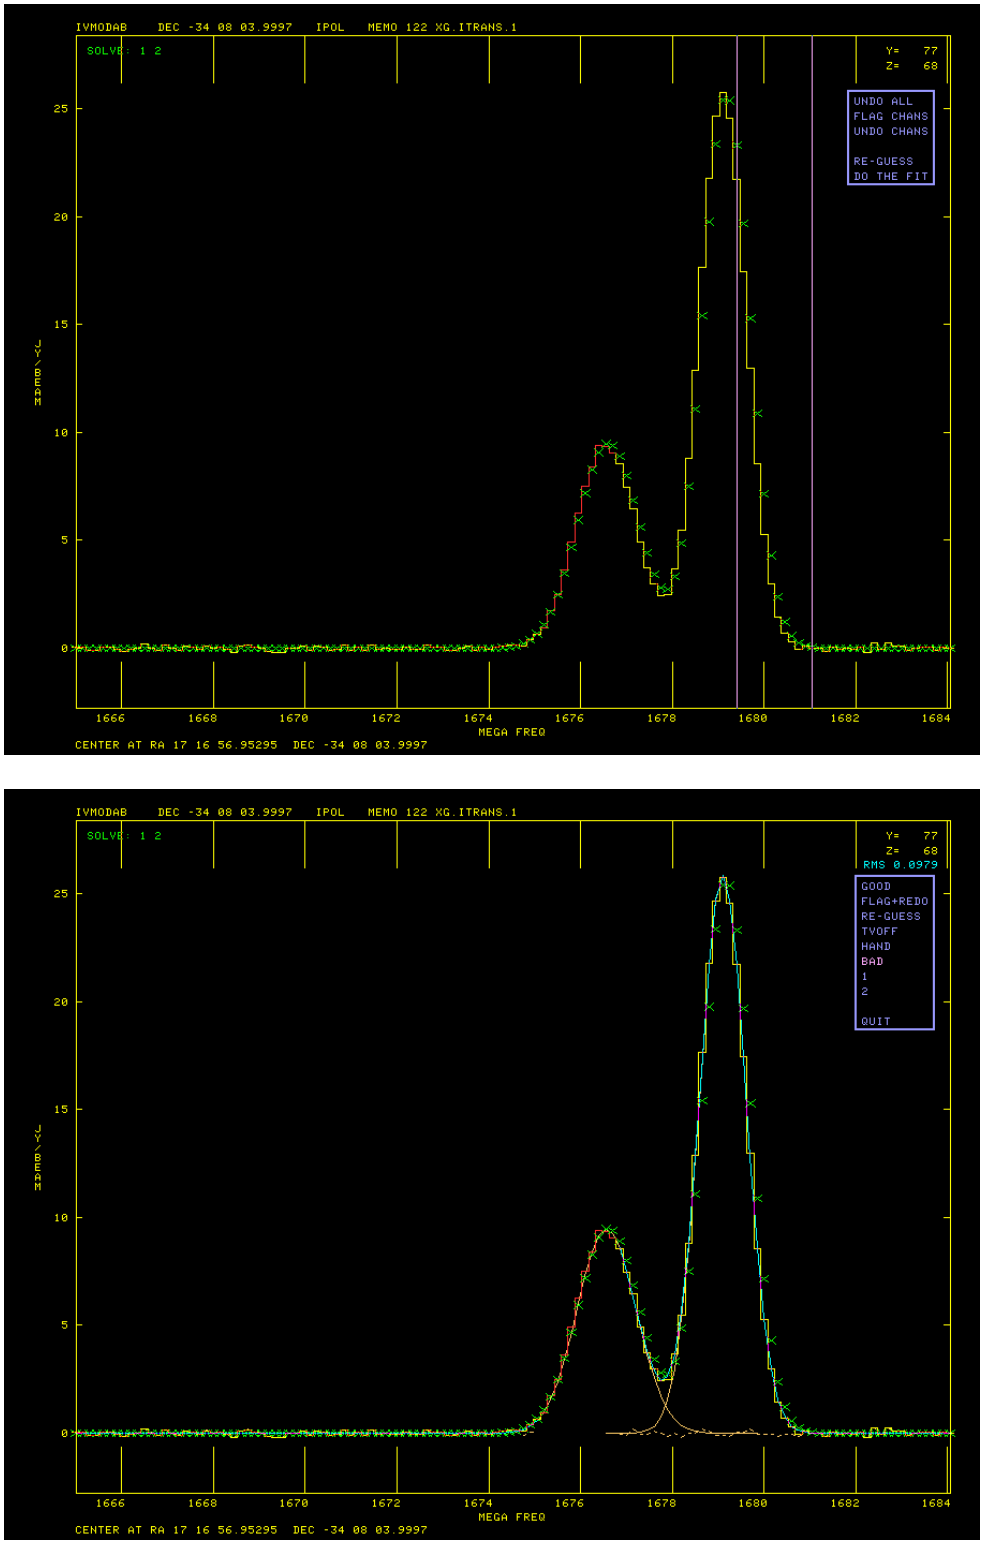

Figure 7: Top:Spectrum to be re-done in flag menu. Some channels are flagged for illustrative (not scientific) purposes and the pink lines show another range of channels which could be flagged. Bottom: Acceptance menu after `FLAG+FIT`. Note that the good initial guess constrains the fitting despite the questionable channel flagging.

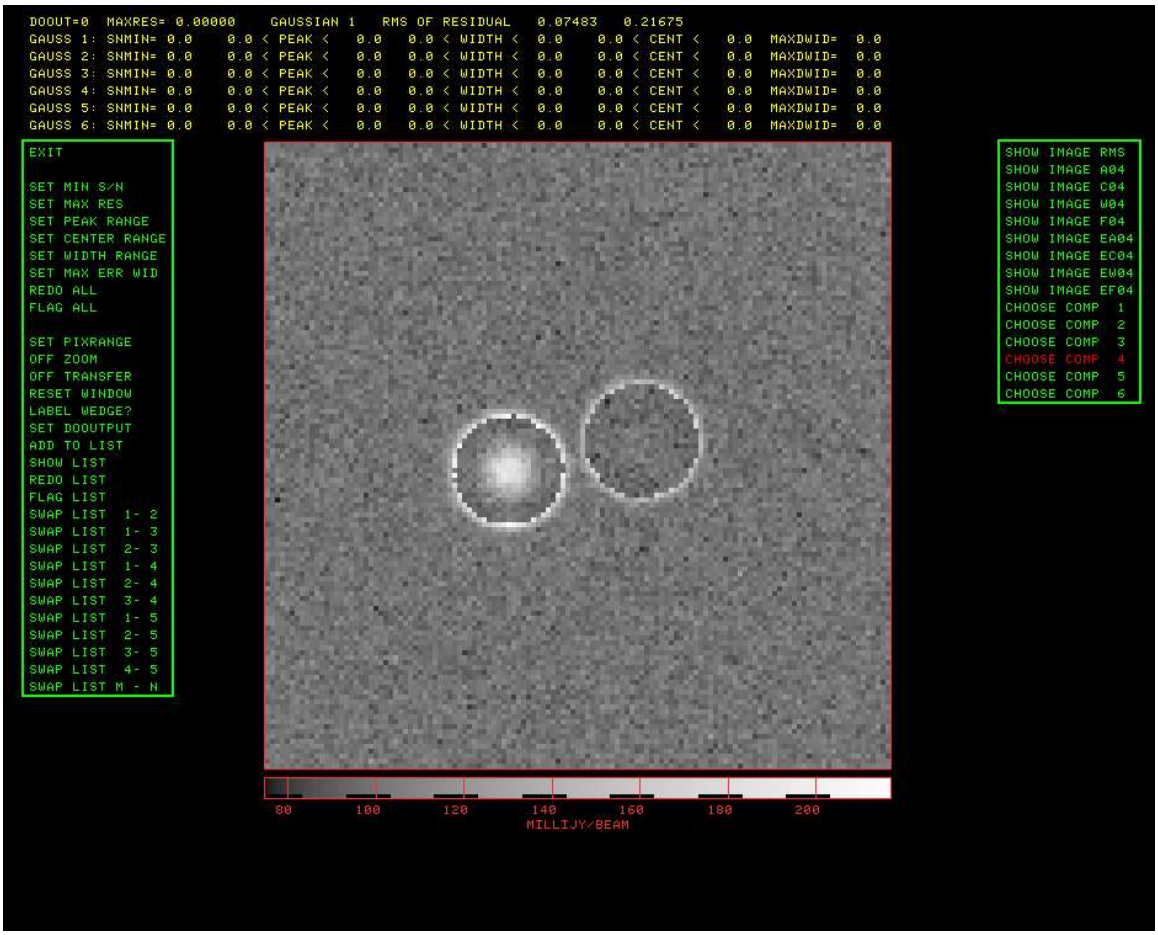

Figure 8: Edit menu when more than 5 components. The rms image is shown with component 4 selected.

SHOW IMAGE RMS	Enter image interaction with the residual rms image
SHOW IMAGE A1	Enter image interaction with peak value of component 1
SHOW IMAGE C1	Enter image interaction with center pixel of component 1
SHOW IMAGE W1	Enter image interaction with width of component 1
SHOW IMAGE F1	Enter image interaction with "flux" of component 1
SHOW IMAGE EA1	Enter image interaction with uncertainty in peak value of component 1
SHOW IMAGE EC1	Enter image interaction with uncertainty in center pixel of component 1
SHOW IMAGE EW1	Enter image interaction with uncertainty in width of component 1
SHOW IMAGE EF1	Enter image interaction with uncertainty in "flux" of component 1

When the XG table is created with more than 5 components, these menus change as illustrated in Figure 8. The SWAP LIST option lists all possible swaps for the first 5 Gaussians. The added option SWAP LIST M - N prompts the user on the terminal for the two components to be switched using the pixel list. The right-hand column changes significantly. Only one component can be viewed at any one time, with the CHOOSE COMP n option to select the desired component. For this figure, the user has looked at the rms image and then selected Gaussian component 4 to be examined. The right-hand menu becomes

SHOW IMAGE RMS	Enter image interaction with the residual rms image
SHOW IMAGE A4	Enter image interaction with peak value of component 4
SHOW IMAGE C4	Enter image interaction with center pixel of component 4
SHOW IMAGE W4	Enter image interaction with width of component 4
SHOW IMAGE F4	Enter image interaction with "flux" of component 4
SHOW IMAGE EA4	Enter image interaction with uncertainty in peak value of component 4
SHOW IMAGE EC4	Enter image interaction with uncertainty in center pixel of component 4
SHOW IMAGE EW4	Enter image interaction with uncertainty in width of component 4
SHOW IMAGE EF4	Enter image interaction with uncertainty in "flux" of component 4
CHOOSE COMP 1	Change attention to Gaussian component 1.
CHOOSE COMP 2	Change attention to Gaussian component 2.
CHOOSE COMP 3	Change attention to Gaussian component 3.
CHOOSE COMP 4	Change attention to Gaussian component 4.
CHOOSE COMP 5	Change attention to Gaussian component 5.
CHOOSE COMP 6	Change attention to Gaussian component 6.

When you select one of the SHOW IMAGE functions most of the following operations will appear in yet another menu. This menu is illustrated in Figure 9 for a one-component image cube and Figure 10 for a 6-component cube. These options mostly invoke familiar functions from *AIPS* to control the FUNCTYPE used in loading the image to the display, to enhance the image intensities, to color the enhanced image intensities, and to zoom the display. Only one of the LOAD AS options will appear, with the next one in the sequence offered when the current one has been invoked. The SET WINDOW option allows you to select a sub-image to view in greater detail, while RESET WINDOW returns to viewing the full image. The SWAP n-m options will appear as needed when NGAUSS > 1 and are used to change your selection of component number n and component m in case XGAUS got confused as in the figure. The NEXT WINDOW option appears when needed to display an image too large to fit on the display screen.

Two operations in this menu are different from the usual. CURVALUE provides the capability of selecting positions for the edit "list." During the CURVALUE operation, position the cursor over the desired pixel and press buttons A, B, or C to add that pixel to the list. The BLOTCH LIST and SWAP n-m operations use a TV blotch operation like that in the *AIPS* verb TVSTAT and task BLSUM. You are to mark with a "blotch" region those pixels in the present image which are to be added to the pixel list or to have their solutions swapped with those of the selected component. Instructions will appear in the message window as you proceed. Begin by positioning the cursor at a pixel to be the first vertex of a connected sequence of vertices and press TV button A. Move to the next vertex and press button A again and repeat until you have marked all vertices for this region. Then press button D if you are done with this region or button C if you need to re-position one of the vertices. In this case, move the cursor to the vertex to be re-positioned, press button A and drag the vertex to the corrected position. Press button A or B to fix that vertex and go on to reset another vertex or D to end this region. You may do this as many times as needed. Note that the swap options only apply between the currently selected and displayed component with all others.

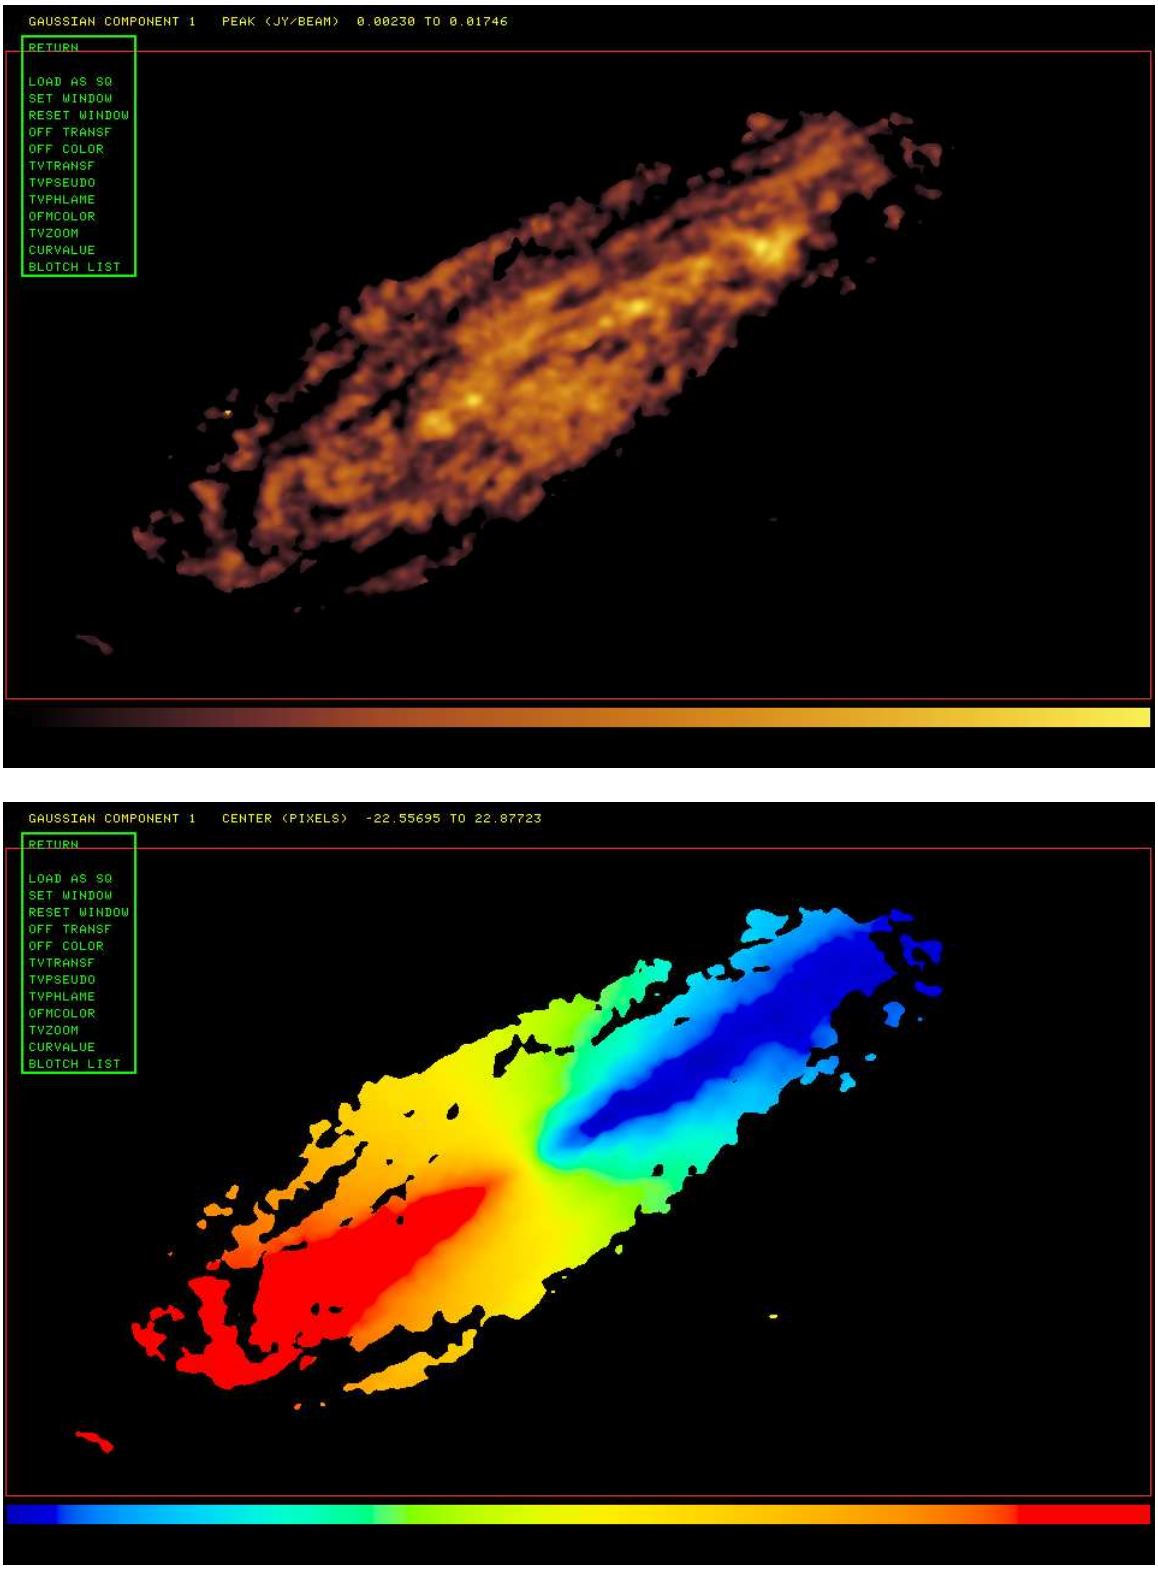

Figure 9: Top: HI galaxy: peak value of one-component fit colored by `OFMCOLOR`; image interaction menu. Bottom: same cube showing the center channel colored by `TVPSEUDO`.

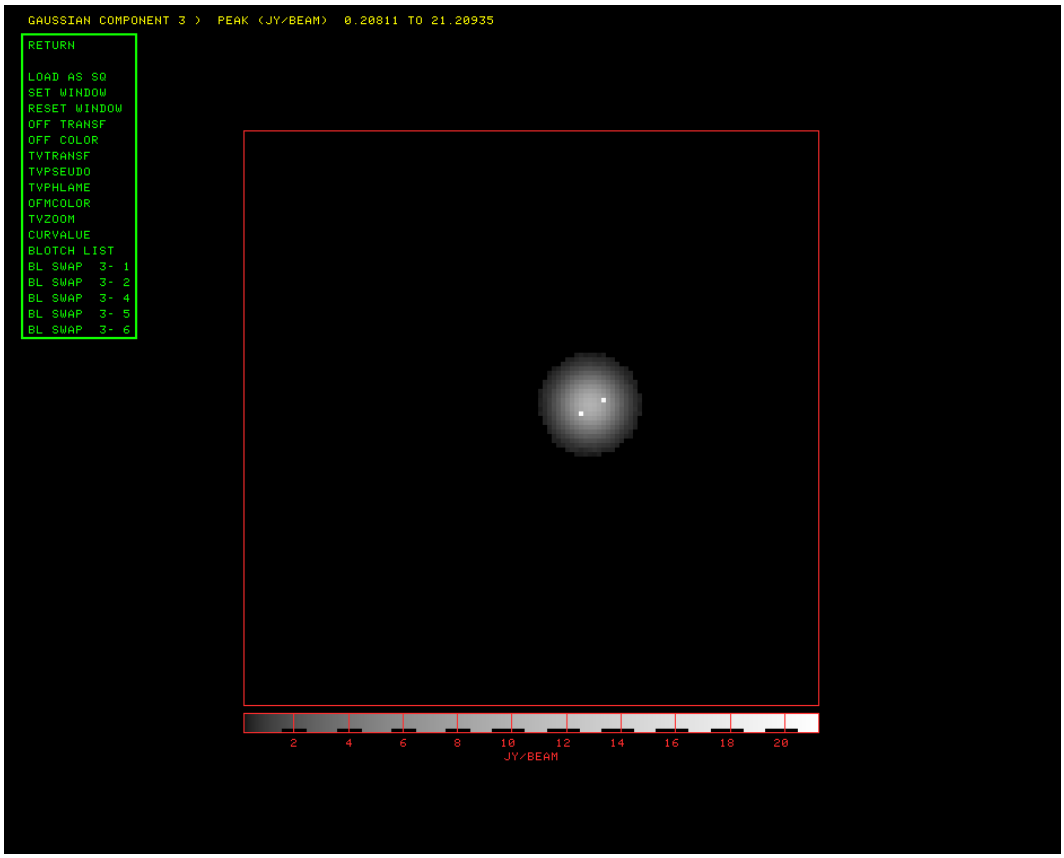

Figure 10: Model source showing menu with 6 components in the XG table.

RETURN	Return to the above menus, image stays displayed
LOAD AS SQ	Re-load image with square root transfer function
LOAD AS LG	Re-load image with log transfer function
LOAD AS L2	Re-load image with extreme log transfer function
LOAD AS LN	Re-load image with linear transfer function
SET WINDOW	Set a sub-image to view
RESET WINDOW	Return too viewing the full image
OFF TRANSF	Turn off enhancement done with TVTRANSF
OFF COLOR	Turn off color enhancements
TVTRANSF	Black & white image enhancement
TVPSEUDO	Color enhancement of numerous sorts
TVPHLAME	Color enhancement of flame type, multiple colors
OFMCOLOR	Many standard color tables (OFMs) available
TVZOOM	Interactive zooming and centering of image
CURVALUE	Display value under cursor, mark pixels for list
BLOTCH LIST	Interactively add a blotch area to the pixel list
SWAP 1-2	Swap solutions for components 1 and 2 interactively
NEXT WINDOW	Move to next window in large images

When you have finished getting the images just the way you want them, you may write them out as AIPS image files. Select the SET DOOUTPUT option until its value, shown at the top of the screen, is what you want. In XGAUS, values 1, 3, 5, and 7 cause a residual image cube to be written, values 2, 3, 6 and 7 cause images of the parameter values and their uncertainties to be written, and values 4, 5, 6, and 7 cause the image of the residual rms to be written. The baseline and slope images and their uncertainties are given class CONST, SLOPE, DCONST, and DSLOPE, while the parameter images and their uncertainties get class AMPLnn, CENTnn, WIDTnn, DAMPnn, DCENnn, and DWIDnn, and the flux and its uncertainty get class FLUXnn, and DFLXnn, where nn is the component number. The residual rms image gets class RESRMS.

3 Zeeman splitting: ZEMAN

When an intrinsically unpolarized spectral line is emitted in the presence of a magnetic field, the right and left circular polarizations have their frequencies shifted in opposite directions by an amount proportional to the magnetic field (at least for modest magnetic fields). The traditional analysis of data to measure this splitting works only for those cases in which the separation of polarizations is a small fraction of the line width. In that case, the function that has been traditional is

$$V(x,y,z) = A(y,z)I(x,y,z) + 0.5B(y,z)\ \frac{dI(x,y,z)}{dx} \tag{2}$$

where $V(x,y,z)$ is the V Stokes polarization component. $I(x,y,z)$ is the unpolarized I Stokes component, x is the spectral axis value, (y,z) is the celestial coordinate value, and $A(y,z)$ and $B(y,z)$ are the parameters to be found by a linear least-squares method.[4] The AIPS task ZEMAN offers this model, with two choices for the method by which the derivative of I with x is computed. If the total intensity spectrum has had XGAUS applied, another function may be used instead:

$$V(x,y,z) = A(y,z)I(x,y,z) + 0.5\sum_{1}^{\text{NGAUSS}} B_i(y,z)\ \frac{dG_i(x,y,z)}{dx} \tag{3}$$

where $G_i(x,y,z)$ is the spectrum of the i'th Gaussian component and one solves for $A(y,z)$ and NGAUSS component frequency separations $B_i(y,z)$. In an early study, VLA data in the C configuration lacked the spatial resolution to separate two source components that emit OH lines at different frequencies. This second method easily found two different magnetic field values. This source has been observed with much better spatial resolution in A configuration which separates the two components spatially. The field values found by ZEMAN matched those found in the published, high-resolution data. Note that ZEMAN returns B in units matching the input x axis (usually Hz or m/sec). Any association with magnetic field values is left to the user.

3.1 Inputs

The inputs for ZEMAN are very similar to those for XGAUS. INNAME *et al.* specify the V polarization cube which must be in transposed form with frequency (or velocity) as the first axis. IN2NAME *et al.* specify the corresponding I polarization cube, similarly transposed. The axes of the two images must match. BLC and TRC define the spectral and celestial coordinate regions of interest for the fit and FLUX gives the lower limit in the I image for the average of three consecutive channels for the spectra to be fit. INVERS specifies the input version of the ZE table in which the ZEMAN results are stored. Zero means to make a new one, other values mean to re-visit an existing solution. Adverb OPTYPE specifies which of the above formulæ are solved. OPTYPE = 'GAUS' says to use the XG file IN2VERS attached to the input I polarization cube to solve using Equation 3. OPTYPE = '2SID' says to solve Equation 2, evaluating the derivative by

$$\frac{dI(x,y,z)}{dx} = 0.5\,[\,I(x+1,y,z) - I(x-1,y,z)\,]$$

[4] Sault, R.J., Killeen, N. E. B.., Zmuidzinas, J., Loushin, R. 1990, *Ap. J.*, **74**, 437-461.

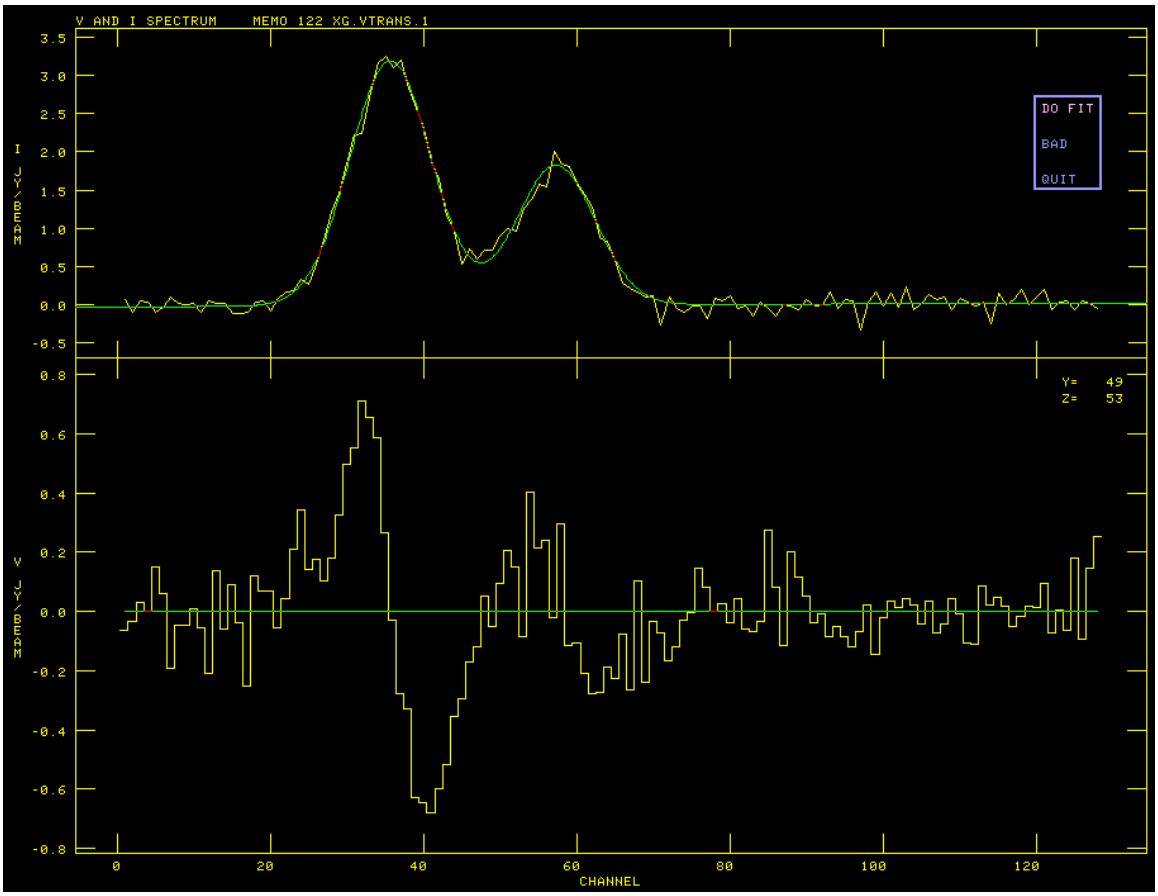

Figure 11: First model spectrum to fit, initial "guess" has all ZEMAN parameters zero.

while any other OPTYPE value says to evaluate the derivative with

$$\frac{dI(x,y,z)}{dx} = I(x,y,z) - I(x-1,y,z)$$

while solving Equation 2. DOOUTPUT controls what files are written — this may be changed interactively (except for one-dimensional images) so leave it zero at this point. Set DOTV> 0 to use TV menus to prompt you. Even though the fitting operation is linear and so needs little or no guidance, it is best to watch what is happening so you should never set this adverb false. If all seems well, you can turn off the TV after watching a few of the solutions. $A(y,z)$ is always fit but should be very close to zero if you are fitting a cube which has already had the leakage term fit and removed. RMSLIMIT is an upper limit for the rms of a fit before the fit is viewed as "failed" which causes the TV and interaction to be turned back on after you have turned it off. You should get a good idea of an appropriate value from your initial uses of ZEMAN or from your knowledge of the noise in your data cube.

Note that ZEMAN does not have an NGAUSS adverb. When OPTYPE = 'GAUS', the number of Gaussians at each pixel and their component number(s) are controlled by the selected XG table. If there is a problem in that table, the user must return to XGAUS to correct it. This is also the reason that the edit stage does not have any "swap" options.

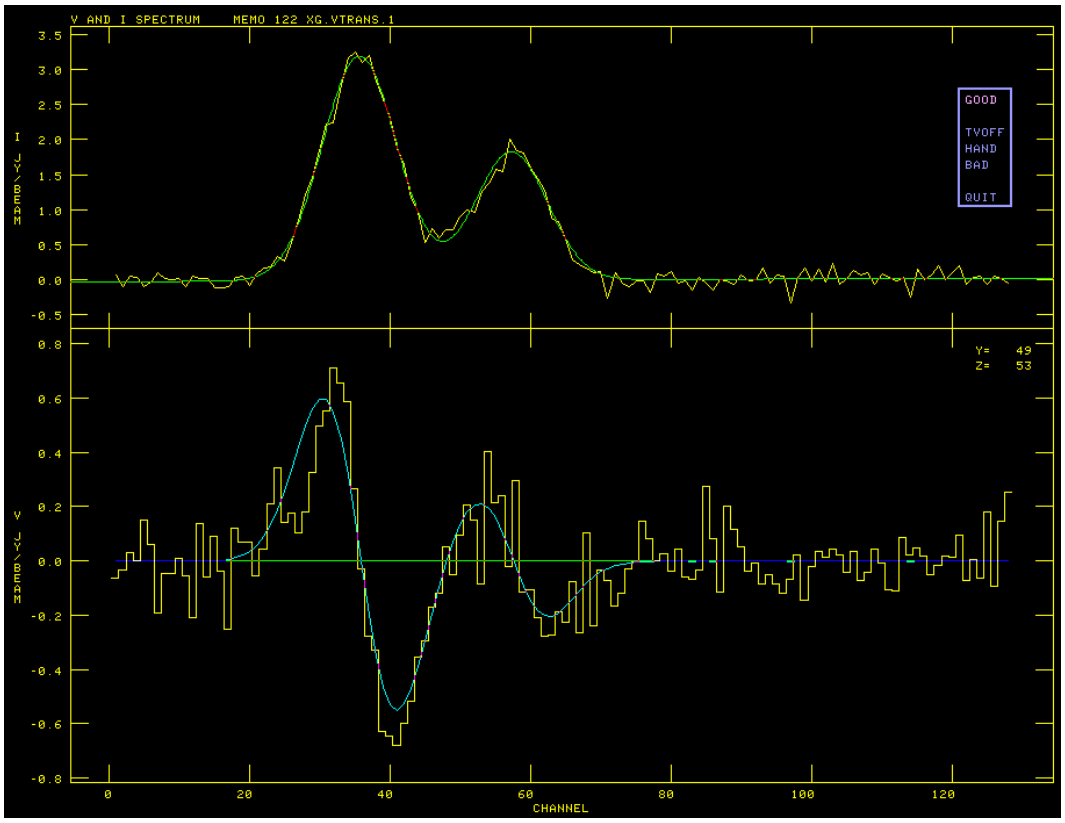

Figure 12: First model spectrum to fit, plot after fitting.

3.2 Fitting

The fitting process starts with a plot of the I polarization spectrum across the top and the V polarization across the bottom. The data and axis labels are plotted in graphics channel one (usually yellow) and the initial guess as lines in graphics channel two (usually green). The Gaussians, when `OPTYPE='GAUS'`, are shown as a smooth line on the I polarization plot, while the initial guess, plotted in the V polarization, is plotted at the locations of the data samples. Then you are offered a menu of options, as shown in the present figures, on the TV (`DOTV` > 0). The first spectrum to be fit (from the model source used for `XGAUS` in this memo) is illustrated in Figure 11 showing that the first guess is zero for all `ZEMAN` parameters. The menu that appears at this point is

`DO FIT`	To do the fitting
`BAD`	To mark this voxel bad and go on to the next
`QUIT`	To exit `ZAMAN` now

Note that there is no `RE-GUESS` since there is no need to enter them with linear least squares. You select menu options in the same way as `XGAUS`, with "button" `D` giving appropriate real-time help information and buttons `A`, `B`, and `C` selecting the highlighted option. The option to `QUIT` causes the task to quit at this point. You may re-start later. The option `BAD` will mark this position as failed and go on to the next position. The option `DO FIT`, currently highlighted, will cause the task to attempt the linear fit with the current initial guess. The result of that fit is shown in Figure 12.

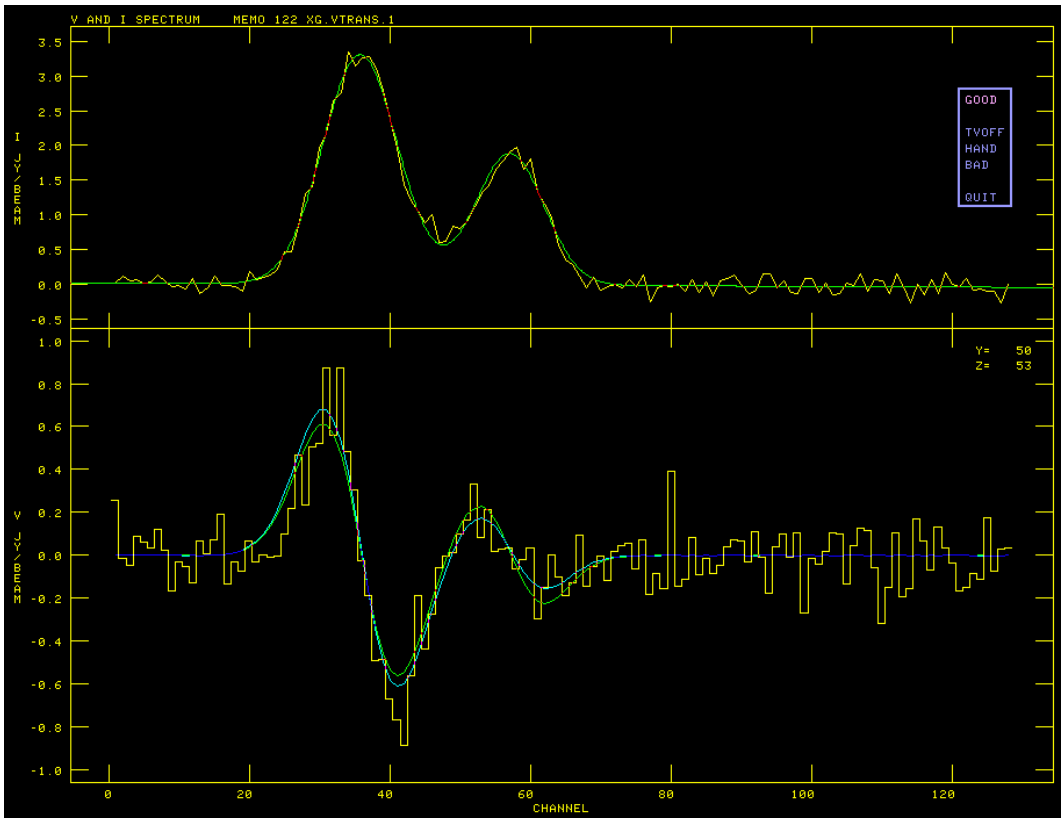

Figure 13: Later model spectrum to fit, initial "guess" in green has Gaussians and values for $B_i(y,z)$, cyan shows actual fit.

The menu after each fit contains

GOOD	To accept the solution and go on
TVOFF	To accept the solution and go on with no interaction
HAND	To enter a new solution from the terminal
BAD	To mark this solution as bad and go on
QUIT	To exit ZAMAN now

What most of these do is obvious. HAND prompts you to enter the Gain value and then either *Enter avg Field in channels* or *Enter Gaussian xx Field in channels* for $xx = 1$ to NGAUSS depending on the OPTYPE selected. When you have entered all values, ZEMAN repeats the plot of Figure 13 with the new values. When the task goes on to the next voxel it starts with the $B_i(y,z)$ of the last fit (Figure 13) and so has a much better initial guess.

3.3 Editing and output

Eventually all positions selected by FLUX and BLC and TRC will have been fit. At this point, the task computes images of the residual spectral rms and of the fit parameters $A(y,z)$ and $B_i(y,z)$ and their uncertainties. This "edit" stage is omitted if the input image is one dimensional, but, in that case, channel flagging is available in the initial fitting of the spectrum. It then offers a menu of options which will allow you to view these images and revisit positions that seem to have produced incorrect fits. This "edit" menu is illustrated

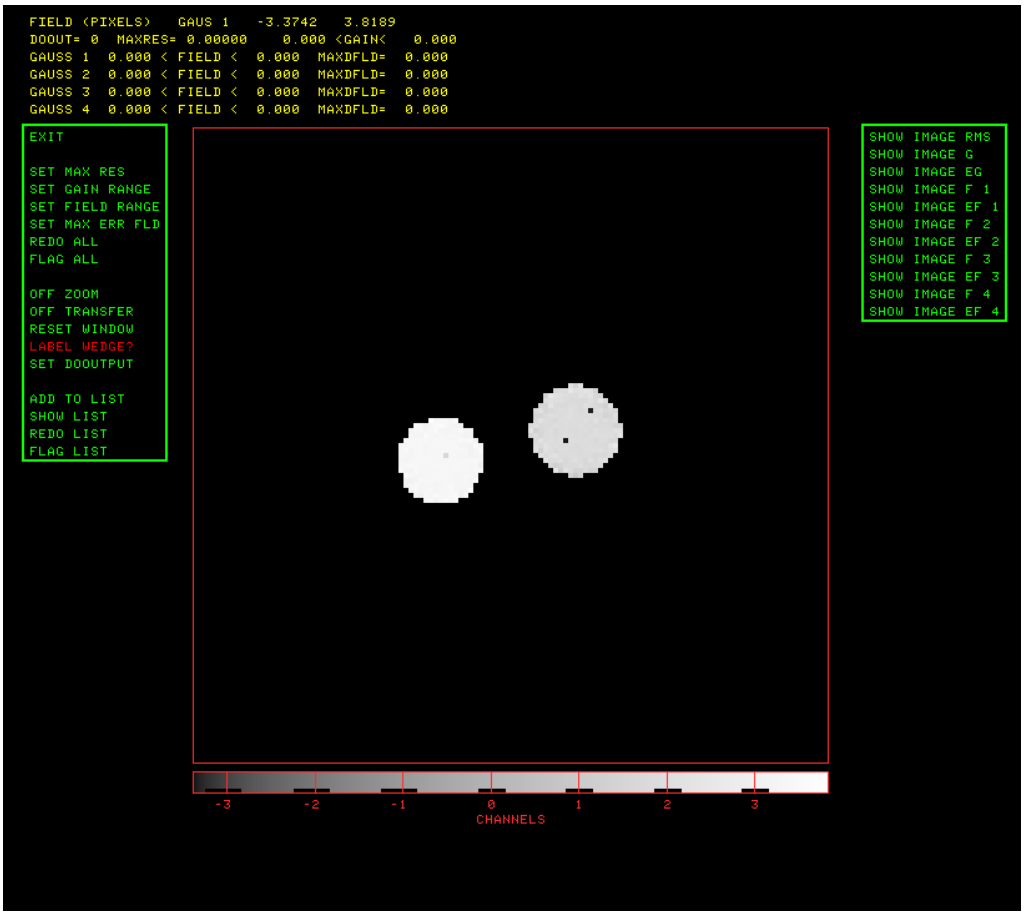

Figure 14: ZEMAN fit (component 1 splitting "field") and image editing screen. Note bad pixel due to bad pixel in XGAUS fits; see Figure 5.

in Figure 14. Note that the size of the region of interest is quite small as illustrated here. If SET WINDOW is used to select a small sub-image, then ZEMAN will blow it up by pixel replication to a reasonable size as in the following figure. The first column of the menu includes:

EXIT	Exit XGAUS, writing output images if DOOUTPUT is now > 0.
SET MAX RES	Set maximum residual for okay solutions
SET GAIN RANGE	Set gain value range(s) for okay solutions
SET FIELD RANGE	Set field range(s) for okay solutions
SET MAX ERR FLD	Set maximum error(s) in field for okay solutions
REDO ALL	Re-do all solutions which are not okay
FLAG ALL	Mark bad all solutions which are not okay
OFF ZOOM	Turn of TV zoom
OFF TRANSFER	Turn off black & white and color TV enhancements
RESET WINDOW	Display full view of current image
LABEL WEDGE?	Turn on/off labeling of step wedge
SET DOOUTPUT	Increment DOOUTPUT in loop $0-3$
ADD TO LIST	Type in output pixel coordinates to add to list
SHOW LIST	Display coordinates in list
REDO LIST	Re-do solutions for all pixels in list
FLAG LIST	Flag solutions for all pixels in list

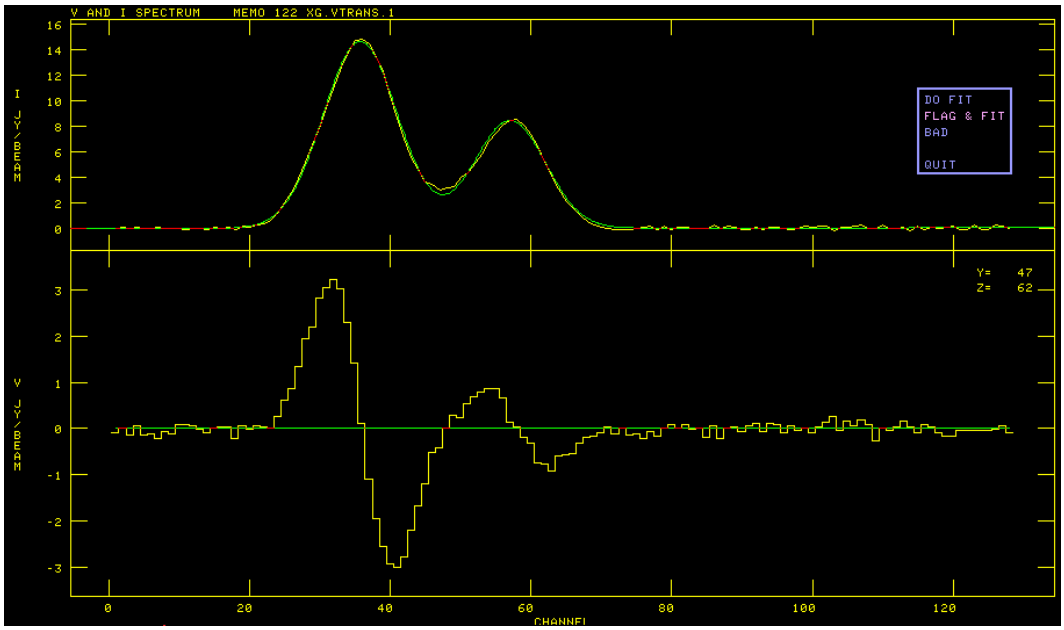

Figure 15: ZEMAN initial display to re-do a fit.

There are two kinds of editing implemented here. In the first, the user establishes the parameter extrema which should be viewed as acceptable. The extrema currently set are shown in the title lines and include the maximum rms, the range of allowed values for "gain" ($A(y,z)$), the range of allowed values for "field" ($B(y,z)$ in pixels), and the maximum uncertainty in the field. If the number of Gaussians is greater than 8, the display of allowed ranges includes only the first eight. Then, ZEMAN may be told to flag all solutions not meeting these criteria, or, more profitably perhaps, to revisit those positions to see why a poor fit was obtained. Note that, unlike the other tasks in this memo, the linear nature of the fit in ZEMAN means that only one solution is possible at each celestial coordinate unless the data are changed in some way. The other editing method is similar, but acts on a list of pixel positions. These may be entered by typing in values, by clicking on suspect pixel positions in the CURVALUE function described below, or by using the BLOTCH LIST function also described below. The contents of the list may be viewed, the solutions at the positions may be flagged, and they may be revisited to see why they are suspect. Note that there are no swapping of solutions in this task; XGAUS establishes which component is which. After the flagging or revisiting, the list is cleared.

In 31DEC25 the operations of REDO ALL and REDO LIST offer an additional menu option.[5] When a spectrum to be re-done is first displayed it is called FLAG & FIT as illustrated in Figure 15. If this option is selected the screen shown in the upper portion of Figure 16 appears. The menu contains

UNDO ALL	Undo all channel flagging
FLAG CHANS	Interactively flag a range of channels
UNDO CHANS	Interactively undo flags in a range of channels
DO THE FIT	Leave the flag routine and do the fit

Flagged channels are plotted in a different color than unflagged channels (usually reddish and yellow, respectively). UNDO ALL simply removes all flags in the current spectrum. FLAG CHANS and UNDO CHANS bring up two vertical lines (usually pink in color). The cursor position controls the position of the lines, button A switches between the two, button B sets a channel range and continues, button C sets a range

[5] This option is available at the beginning when the input image is one-dimensional.

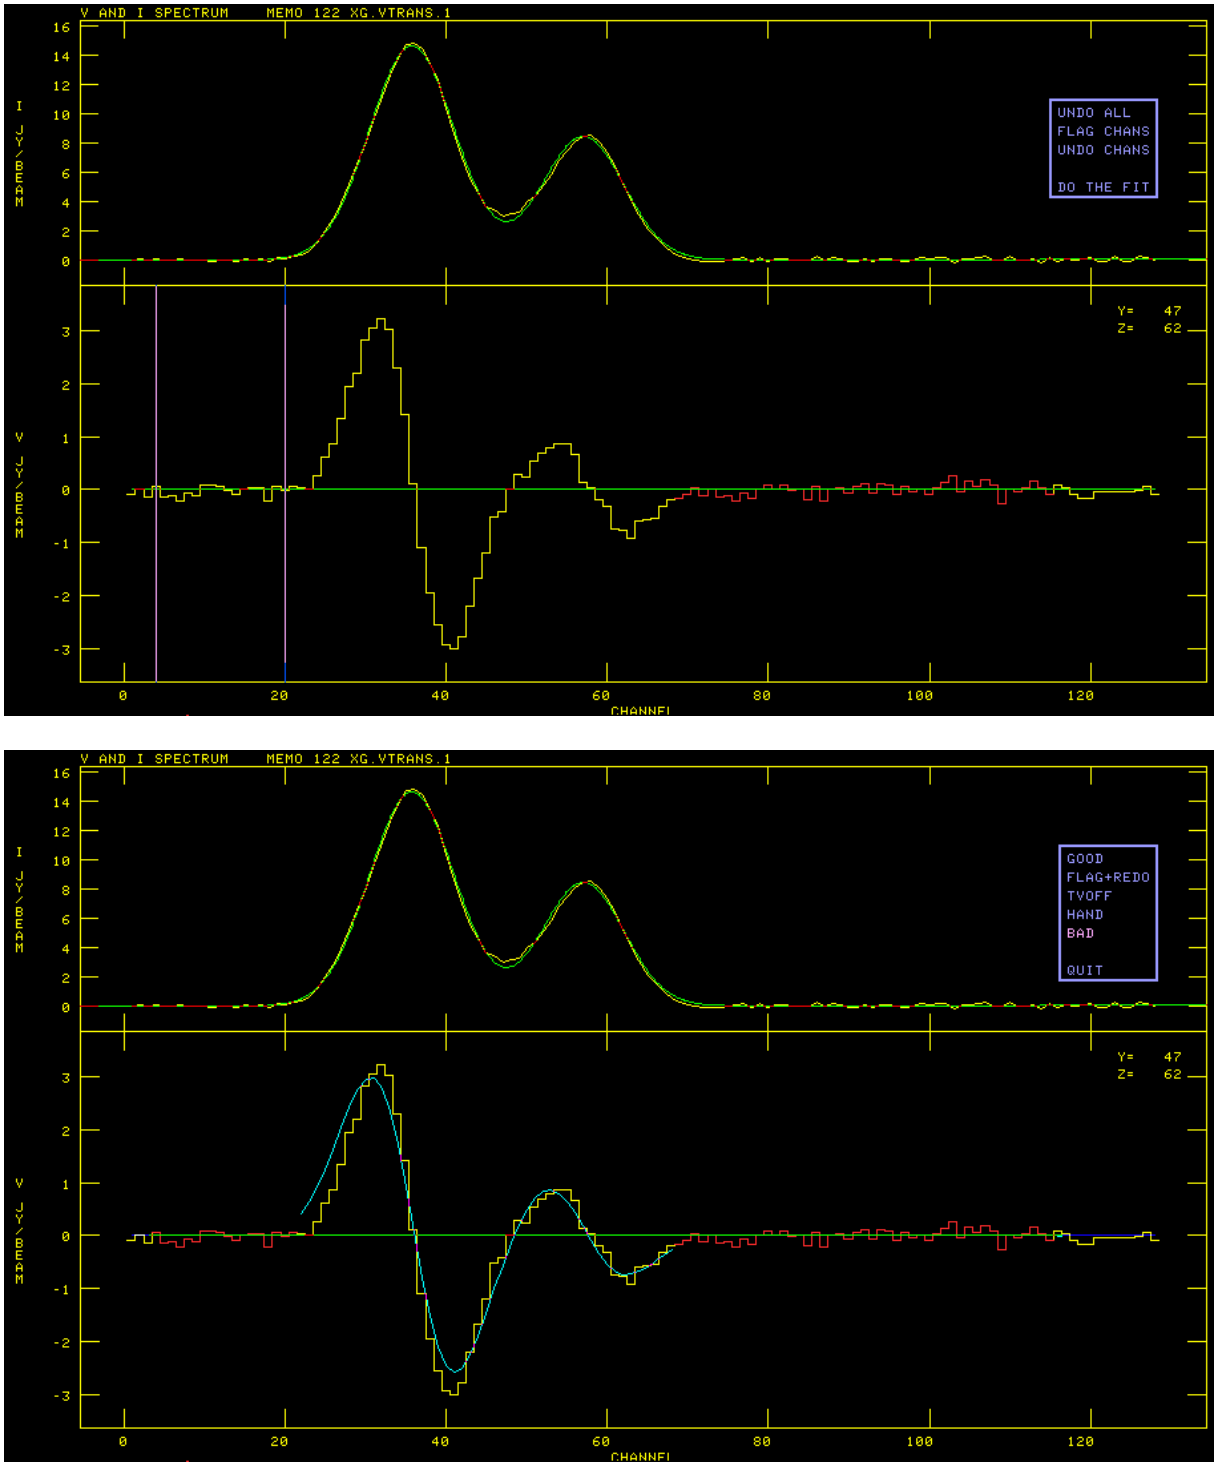

Figure 16: Top: Spectrum to be re-done in flag menu. Some channels on the right hand side are flagged and the pink lines show another range of channels which will also be flagged. Bottom: Acceptance menu after `FLAG & FIT`.

and exits, and button D exits without setting a new range. Finally DO THE FIT attempts the Zeeman fit and returns to the menu to accept the result. That screen is illustrated in the bottom of Figure 16. Note the FLAG+REDO option added to this menu in the edit phase. That option takes you back through the sequence just illustrated.

The second column of the main edit menu contains

SHOW IMAGE G	Enter image interaction with gain
SHOW IMAGE EG	Enter image interaction with uncertainty in the gain
SHOW IMAGE F1	Enter image interaction with field of component 1
SHOW IMAGE EF1	Enter image interaction with uncertainty in field of component 1
SHOW IMAGE F2	Enter image interaction with field of component 2
SHOW IMAGE EF2	Enter image interaction with uncertainty in field of component 2

There is one or, if OPTYPE='GAUS' as in the illustrated cases, the maximum of NGAUSS of the Fn and EFn options. When you select one of the SHOW options, the options in yet another menu appear along with a display of the selected image. As illustrated in Figure 17, these options are

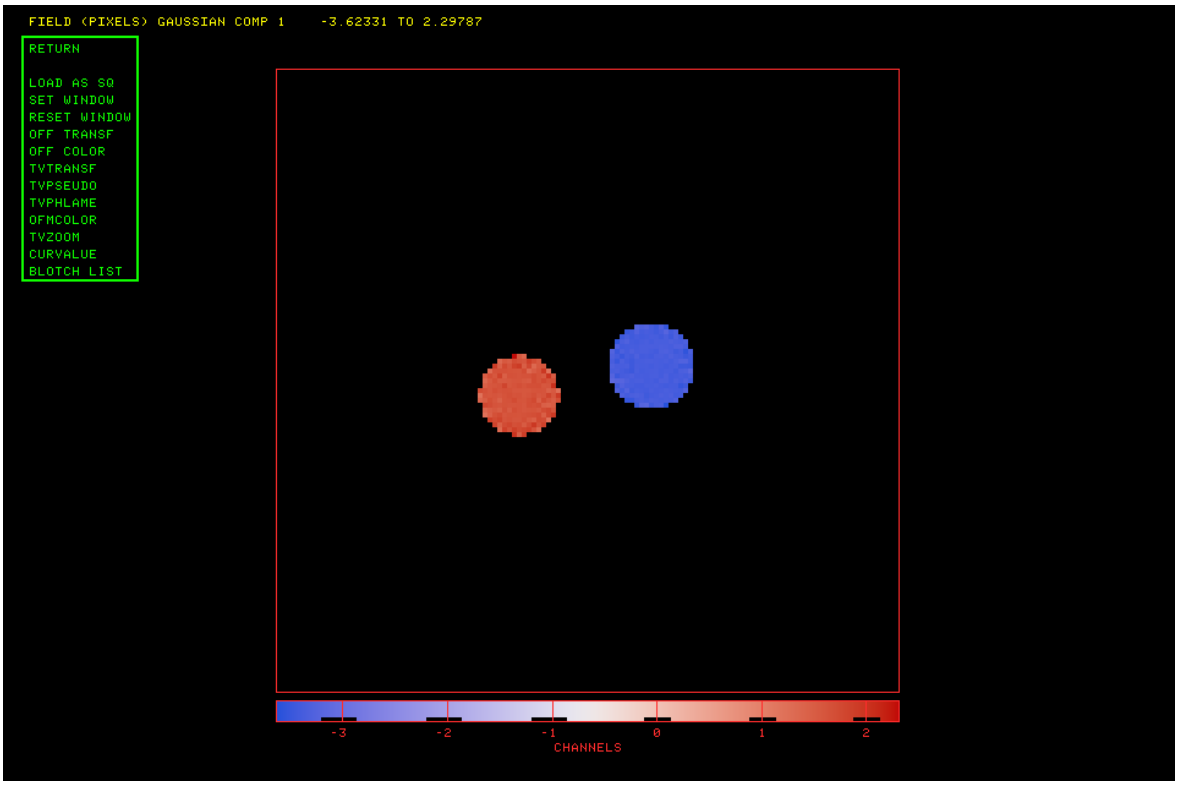

Figure 17: ZEMAN image of component 1 field $B_1(y, z)$ with display options.

RETURN	Return to the above menus, image stays displayed
LOAD AS SQ	Re-load image with square root transfer function
LOAD AS LG	Re-load image with log transfer function
LOAD AS L2	Re-load image with extreme log transfer function
LOAD AS LN	Re-load image with linear transfer function
SET WINDOW	Set a sub-image to view
RESET WINDOW	Return to viewing the full image
OFF TRANSF	Turn off enhancement done with TVTRANSF
OFF COLOR	Turn off color enhancements
TVTRANSF	Black & white image enhancement
TVPSEUDO	Color enhancement of numerous sorts
TVPHLAME	Color enhancement of flame type, multiple colors
OFMCOLOR	Many standard color tables (OFMs) available
TVZOOM	Interactive zooming and centering of image
CURVALUE	Display value under cursor, mark pixels for list
BLOTCH LIST	Interactively add a blotch area to the pixel list
NEXT WINDOW	Move to next window in large images

Only one of the LOAD AS options will appear, namely the next after the current transfer function from the list of linear, square root, log, and more extreme log transfer functions. CURVALUE provides the capability of selecting positions for the edit "list." During the CURVALUE operation position the cursor over the desired pixel and press buttons A, B, or C to add that pixel to the list. BLOTCH LIST is described in detail in the XGAUS and RMFIT sections of this Memo. The other options are familiar as AIPS verbs. Instructions for interaction will appear on the terminal and button D in the menu may be used to obtain a helpful display on the terminal.

When you have finished getting the images just the way you want them, you may write them out as *AIPS* image files. Select the SET DOOUTPUT option until its value, shown at the top of the screen, is what you want. In ZEMAN, values 4, 5, 6, 7, 12, 13, 14, and 15 cause a residual (data-model) image cube to be written, values 2, 3, 6, 7, 10, 11, 14, and 15 cause images of the gain and field values and their uncertainties to be written, values 1, 3, 5, 7, 9, 11, 13 and 15 cause an image of the V cube to be written with the gain times the I-polarization cube subtracted, and finally values greater than 7 cause an image of the residual rms to be written. The corrected V-polarization image gets the class specified by OUTCLASS, the residual image gets class VRESID, the gain and its uncertainty get classes GAIN and DGAIN, and the field(s) and their uncertainties get classes FELDnn and DFELnn, with $nn = 1$ to the maximum NGAUSS. The residual rms image gets class RESRMS.

4 Polarization fitting: RMFIT

Faraday rotation occurs when a polarized source of radiation $F_0(\Phi_1)$ has its radiation pass through a region of electrons (a "Faraday screen"). The complex polarization $(Q + iU)$ is rotated in position angle by an amount proportional to the wavelength squared, *i.e.*, it is multiplied by $\exp(i2\phi_1\lambda^2)$. If this radiation is then joined by another source of radiation $F_1(\Phi_2)$ and the two then pass through additional screens ϕ_2, ϕ_3, and ϕ_4 after which there is a third source of radiation and one final screen, the net polarization observed is then

$$P(\lambda^2) = F_0(\Phi_1)\, e^{i2\Phi_1\lambda^2} + F_1(\Phi_2)\, e^{i2\Phi_2\lambda^2} + F_2(\Phi_3)\, e^{i2\Phi_3\lambda^2}, \quad (4)$$

where

$$\begin{aligned} \Phi_1 &= \phi_1 + \phi_2 + \phi_3 + \phi_4 + \phi_5 \\ \Phi_2 &= \phi_2 + \phi_3 + \phi_4 + \phi_5 \\ \Phi_3 &= \phi_5 \end{aligned}$$

or, in the limit:

$$P(\lambda^2) = \int F(\Phi)\, e^{i2\Phi\lambda^2} d\Phi \quad (5)$$

"Faraday Rotation Measure Synthesis" is then the Fourier analysis of Equation 5 as initially described by Brentjens and de Bruyn.[6] *AIPS* task FARS takes as input two matching spectral cubes, one in Q and one in U polarization. It does a suitable Fourier transform in wavelength squared and can do a one-dimensional complex Clean and restoration at each spatial coordinate. A sample of a two-component model is shown in Figure 18. Although the model contains data from frequencies from 1 GHz to 3 GHz, the figure shows that it is not possible with this analysis to separate the two components even though they are at a physically realistic and interesting separation.

Therefore a new fitting task, called RMFIT, was written. Its job is to do a non-linear fit of

$$Q(i) = \sum P_j \cos(2\theta_j + 2R_j\lambda(i)^2) F(\beta_j, , \lambda(i)^2) \quad (6)$$

$$U(i) = \sum P_j \sin(2\theta_j + 2R_j\lambda(i)^2) F(\beta_j, , \lambda(i)^2) \quad (7)$$

where P_j is the polarized flux of component j, θ_j is the polarization angle of component j, R_j is the rotation measure of component j, $\lambda(i)$ is the wavelength of spectral channel i, β_j is a "thickness" parameter and F is one of

$$\begin{aligned} F(\beta_j, \lambda(i)^2) &= \left(\frac{\lambda_1^2}{\lambda(i)^2}\right)^{\frac{\beta_j}{2}} \\ &= \frac{\sin(\beta_j\lambda(i)^2)}{\beta_j\lambda(i)^2} \end{aligned}$$

[6] Brentjens, M. A. and de Bruyn, A. G. 2005, *Astron. & Ap.*, 1217.

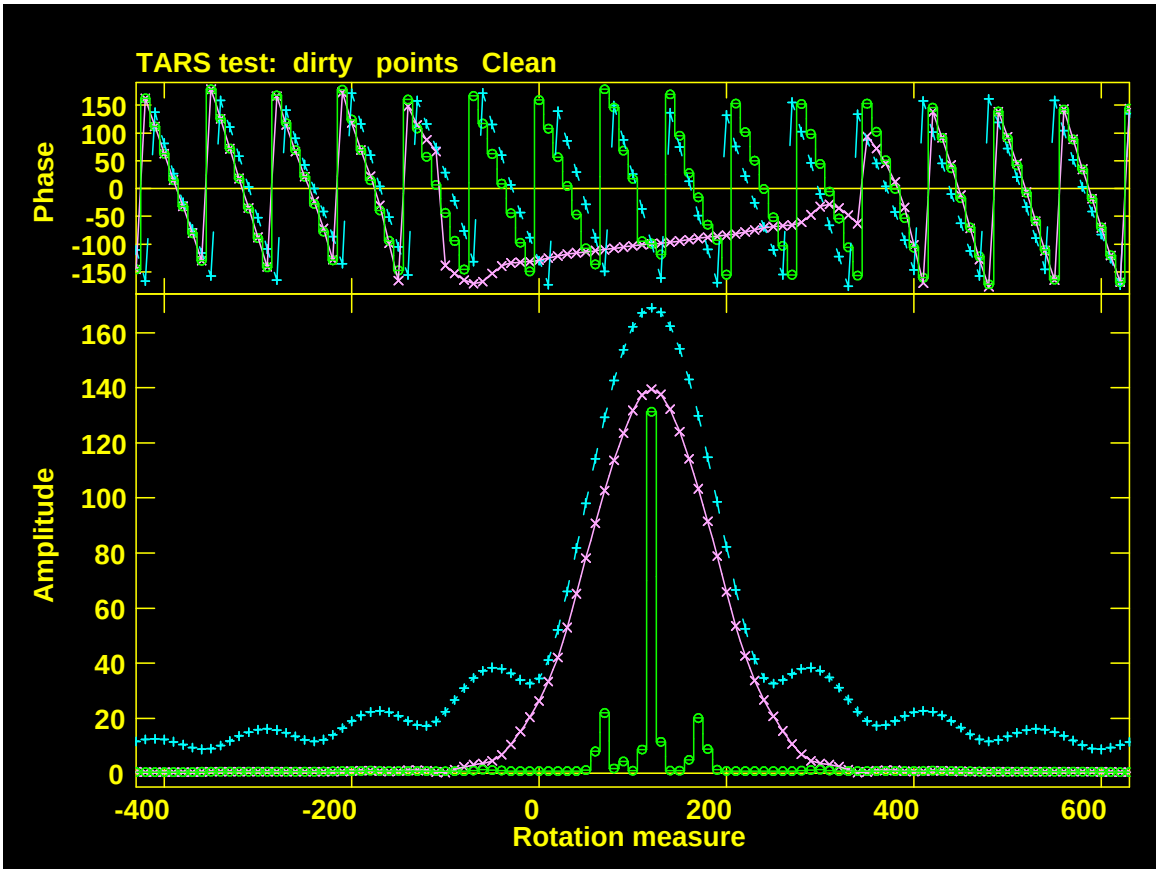

Figure 18: FARS output with model data. Source has components at 100 and 140 radians per m^2. Blue is without Clean, green is Clean restored with very narrow components, and pink is Clean with a normal restoration. Note that Clean finds components at 120 plus about 68 and 172 radians per m^2.

$$
\begin{aligned}
&= e^{-\ln(2)\left(\frac{\beta_j \lambda(i)^2}{1.8954}\right)^2} \\
&= e^{-\ln(2)\frac{\beta_j \lambda(i)^2}{1.8954}}
\end{aligned}
$$

where λ_1 is the wavelength at the fiducial frequency of 1 GHz and the thickness models will be referred to by the names spectral index, slab, Gaussian, and exponential, respectively. The peculiar number 1.8954 was chosen so that the last three models reach the 0.5 point at the same value of β_j. Rotation measure thickness, in general, makes the polarized signal decrease at longer wavelengths. It is one of the reasons that radio sources at low frequency tend to have little or no polarized flux density.

4.1 Inputs

RMFIT expects five input cubes which contain, in order, the Q polarization, the U polarization, the I polarization, the FARS output real or amplitude, and the FARS output imaginary or phase images. All five must align appropriately, the first three must be transposed so that frequency (or frequency ID number) is the first axis (as they go into FARS), and the last two have the rotation measure axis first (as they come out of FARS). The FARS images are required to allow the user to make initial guesses for the model. Do some Cleaning in FARS and restore with a "Clean beam" reasonably close to the default. The I polarization image is optional and may be a continuum image rather than a spectral cube.

The rest of the inputs to RMFIT are much like those of XGAUS. INVERS defines the version of the RM table to

be used to hold the model fit, with zero causing a new table to be used. BLC and TRC define the pixel ranges to be used in the current execution, where BLC(1) and TRC(1) control the spectral channels that will be fit and BLC(2), TRC(2), BLC(3), and TRC(3) control the area in celestial coordinates over which the fitting is done. BLC(1) and TRC(1) do not apply to the FARS images, but they do control which channels are used to find the peak polarization and unpolarized brightnesses for new RM tables. The RM table now contains all pixels on axes 2 and 3 of the input images, so there should never be a problem resuming fitting with a new region of interest. Adverbs YINC and ZINC control the stride taken in the first pass through the cube; a second pass will then fit all pixels not fit in the first pass. Adverbs PCUT and ICUT control which pixels are fit. Any spectra having average total polarization ($\sqrt{Q^2+U^2}$) greater than PCUT and, if an I image is given, total intensity greater than ICUT will be fit.

DOOUTPUT controls what files are written — this may be changed interactively so leave it zero at this point. DOTV is *required* to be true and will use TV menus to prompt you. The badly named adverb NGAUSS controls the maximum number of components to fit and is limited to 4.[7] Adverb DOSPIX controls the choice of thickness model to be fit, with 1 for spectral index, 2 for slab, 3 for Gaussian, 4 for exponential, and anything else for no thickness. Polarization cubes may have quite different noise levels in different spectral channels. Adverb DOWEIGHT controls whether a weighted fit is done. The weights may be found by robust rms methods from the Q and U images or read from a text file specified by adverb INFILE. RMSLIMIT is an upper limit for the rms of a fit before the fit is viewed as "failed" which causes the TV and interaction to be turned back on after you have turned it off. You should get a good idea of an appropriate value from your initial uses of RMFIT or from your knowledge of the noise in your data cube.

A recommended strategy, when creating an RM table, would be to set the the spectral region BLC(1) through TRC(1) to encompass all reliable spectral channels so that the full range of polarized and unpolarized brightnesses will be found. The task begins by creating this table, including the entire input image, and populating each row with the average unpolarized and polarized brightnesses in each spectrum. RMFIT reads the table every YINC rows and ZINC planes and, for those with brightnesses exceeding PCUT and ICUT, attempts a fit. Your interaction with this fit will be described below. After the first pass, the task loops over every row and plane fitting those positions which have enough brightness and which have not already been fit. Finally, after all pixels above the cutoffs have been fit, the task goes into an "edit" mode. It constructs images of each fit parameter and Q and U at zero wavelength and of the uncertainties in these parameters. You may view these images, select positions explicitly or by their parameter values or rms and revisit the fits of those positions. This stage will be described in detail below.

At any time you may exit the task and then re-start it using the same RM table. Good reasons for doing this include fitting smaller regions with each pass using the appropriate number of components for that region. Doing small regions which will have similar parameter values helps a great deal with the initial guessing done by the task (mostly using the previous solution). You might also fit the cube initially with a high value of PCUT and then re-start with a lower value to extend the areas fit.

4.2 Fitting

The fitting process starts with a plot of the FARS amplitude and phase as a function of rotation measure at the first coordinate with adequate signal. The data and axis labels are plotted in graphics channel one (usually yellow) and the initial guess as X's in graphics channel two (usually green). Then you are offered a menu of options, as shown in the present figures, on the TV. The first spectrum to be fit is illustrated in Figure 19 showing that the first guess for one component is sensible. The menu that appears at this point is

[7] Unlike XGAUS and ZEMAN the RM table is always written to allow the maximum number of components (4).

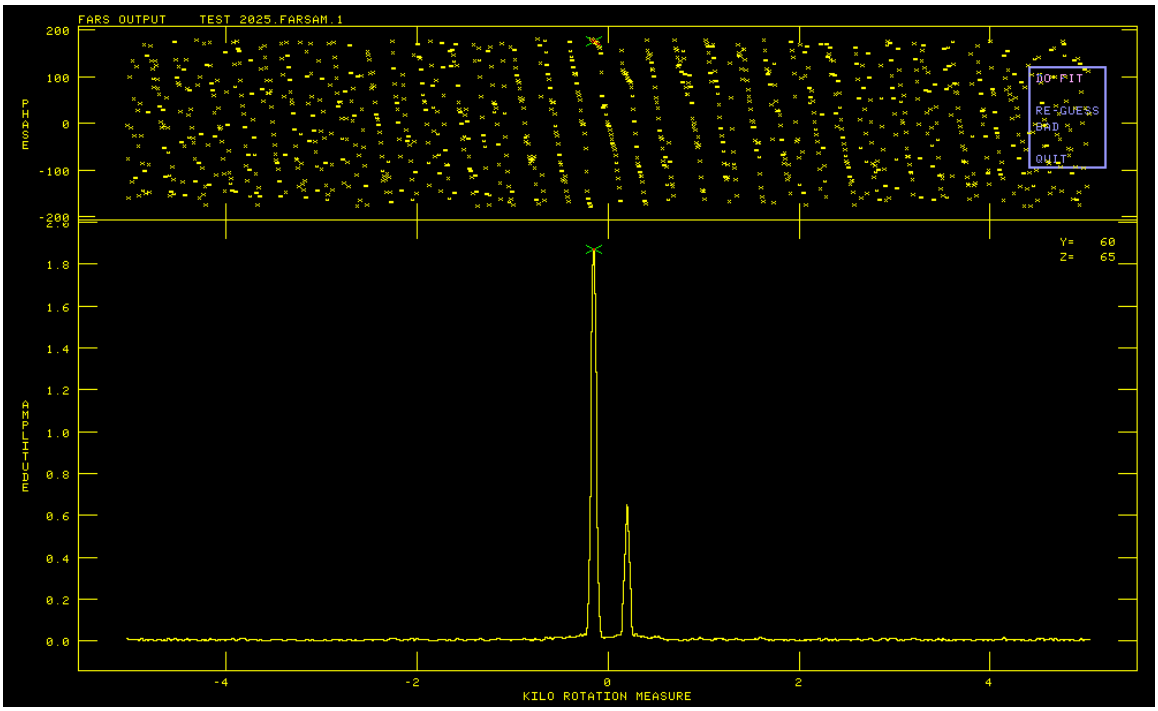

Figure 19: Model field first pixel FARS rotation measure spectrum to fit. With NGAUSS=1 the initial guess is pretty good, but it is clear that this position has 2 components.

DO FIT	RMFIT will attempt the fit
RE-GUESS	RMFIT will prompt you to point out the RM values of the components
BAD	RMFIT will mark this fit as failed and go on
QUIT	RMFIT will exit now

You select a menu option by moving the cursor to the desired option with the mouse and registering that move with the TV by clicking the left mouse button. The selected menu item will change color as shown in the figure. If you press TV "button" D at this point (actually keyboard character D), helpful information about the selected item will appear on your terminal window. If you press one of TV "buttons" A, B, or C (actually keyboard characters A, B, C), the selected function will be performed. The option to QUIT causes the task to quit at this point. You may re-start later. The option BAD will mark this position as failed and go on to the next position. The option DO FIT will cause the task to attempt the non-linear fit with the current initial guess. The selected option in Figure 19 is DO FIT since, for one component, the guess is as good as it can be.

After a fit has been attempted, the U spectrum is plotted in the top half of the screen and the Q spectrum is plotted in the lower half. If you have selected SHOW A&P previously, then the polarization angle spectrum is plotted in the top half and the total polarization spectrum is plotted in the lower half. The data and labels are in graphics channel 1 (usually yellow), the initial guess is plotted in channel 2 (usually green), and the fit model is plotted in channel 4 (usually cyan). The fit values and rmses are displayed on your terminal. The user chose to fit only one component initially, getting the unhappy result shown in Figure 20.

The user, having found that there are actually two components, quit and then restarted RMFIT setting NGAUSS=2 and INVERS to point at the new RM file. Figure 21 shows the initial RM spectrum over again, but with two components the automatic initial guess is not useful. Selecting RE-GUESS, the display in the bottom of Figure 21 appears in a different color to signify that an initial guess is requested. Position the cursor at the peak of each component and press any TV button. The resulting fit is much better as seen in Figure 22.

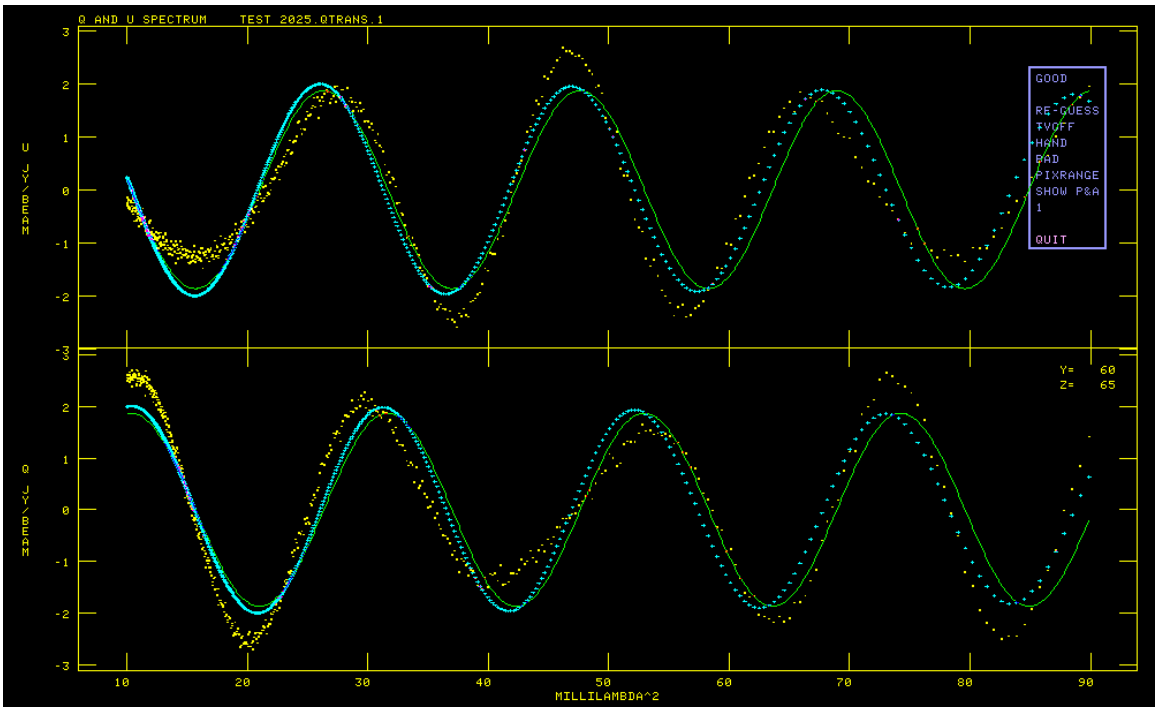

Figure 20: Model field first pixel Q and U spectra, fit with only one component. Initial guess (green) and fit (cyan) are close but miss the data (yellow) badly. User selects QUIT to restart with 2 components.

The menu that appears at this point offers the option to QUIT which causes the task to quit at this point. You may re-start later. The option GOOD will save this solution and go on to the next position. The option BAD will mark this position as failed and go on to the next position. The option RE-GUESS will loop back to prompt you for a new guess and repeat the fit. Options 1, 2, . . ., NGAUSS will loop back to request an initial guess with the selected number of Gaussians. Option HAND will prompt you to enter using the terminal the parameters for each component. Enter on one line for each component, the polarization brightness (in image units), the polarization angle (in degrees), the rotation measure and, if you are fitting a thickness, the thickness width (in radians per meter squared). RMFIT will then repeat the display in Figure 22 to see if you made a good guess. Immediately after a HAND operation only, the option DO FIT is offered to go back with the hand-entered values as the initial guess for a new fit.

GOOD	RMFIT saves the result and goes on
DO FIT	RMFIT re-does the fit
RE-GUESS	RMFIT requests a new initial guess
TVOFF	RMFIT saves the result and turns off interactivity
HAND	RMFIT requests a new guess with the terminal
BAD	RMFIT marks result bad and goes on
SHOW Q&U	RMFIT switches plot to Q and U
SHOW P&A	RMFIT switches plot to amplitude and phase
1	RMFIT requests a new initial guess with one component
2	RMFIT requests a new initial guess with two components
QUIT	RMFIT exits immediately

The user may select one of the SHOW options to switch between displays of Q and U and displays of total polarization and angle as in Figure 22. Option TVOFF allows you to turn off interactivity, allowing the task to run using its own initial guesses until it finds a completely unreasonable solution or one with an rms greater

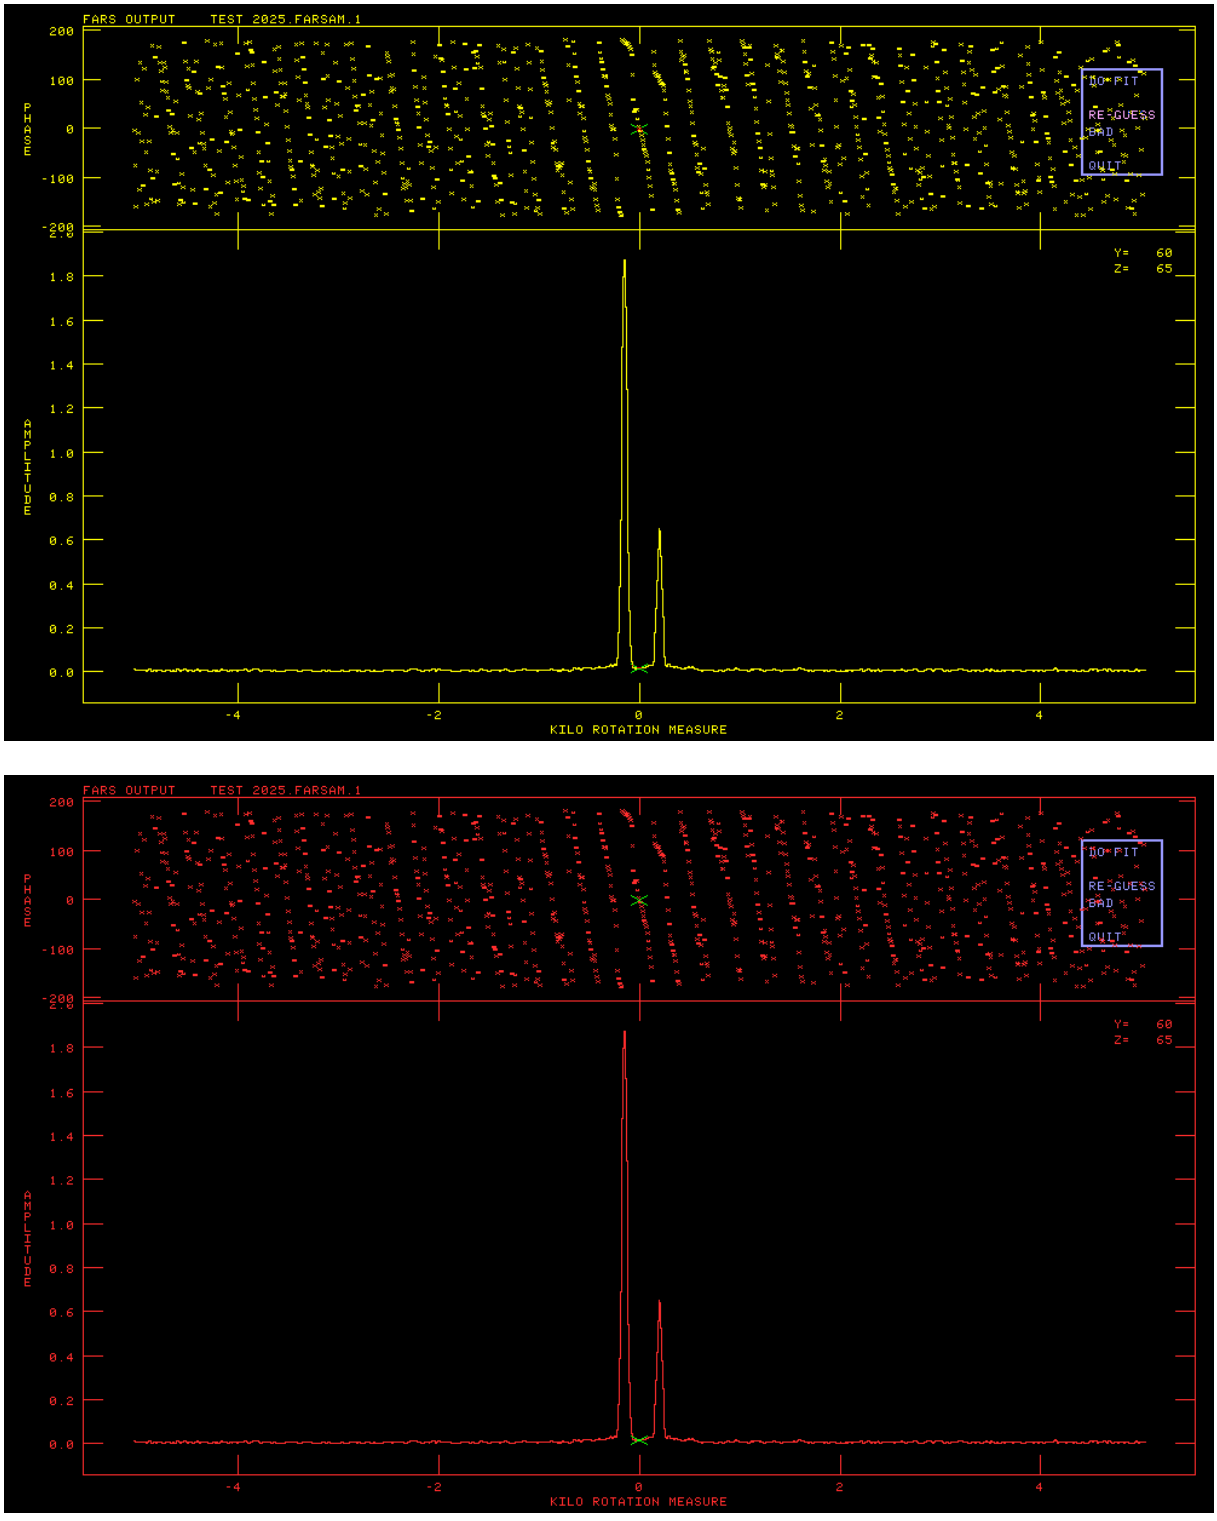

Figure 21: Top: RM spectrum with `NGAUSS = 2`. Initial guess is not good so `RE-GUESS` is selected. Bottom: Display when entering a new initial guess is in a different color.

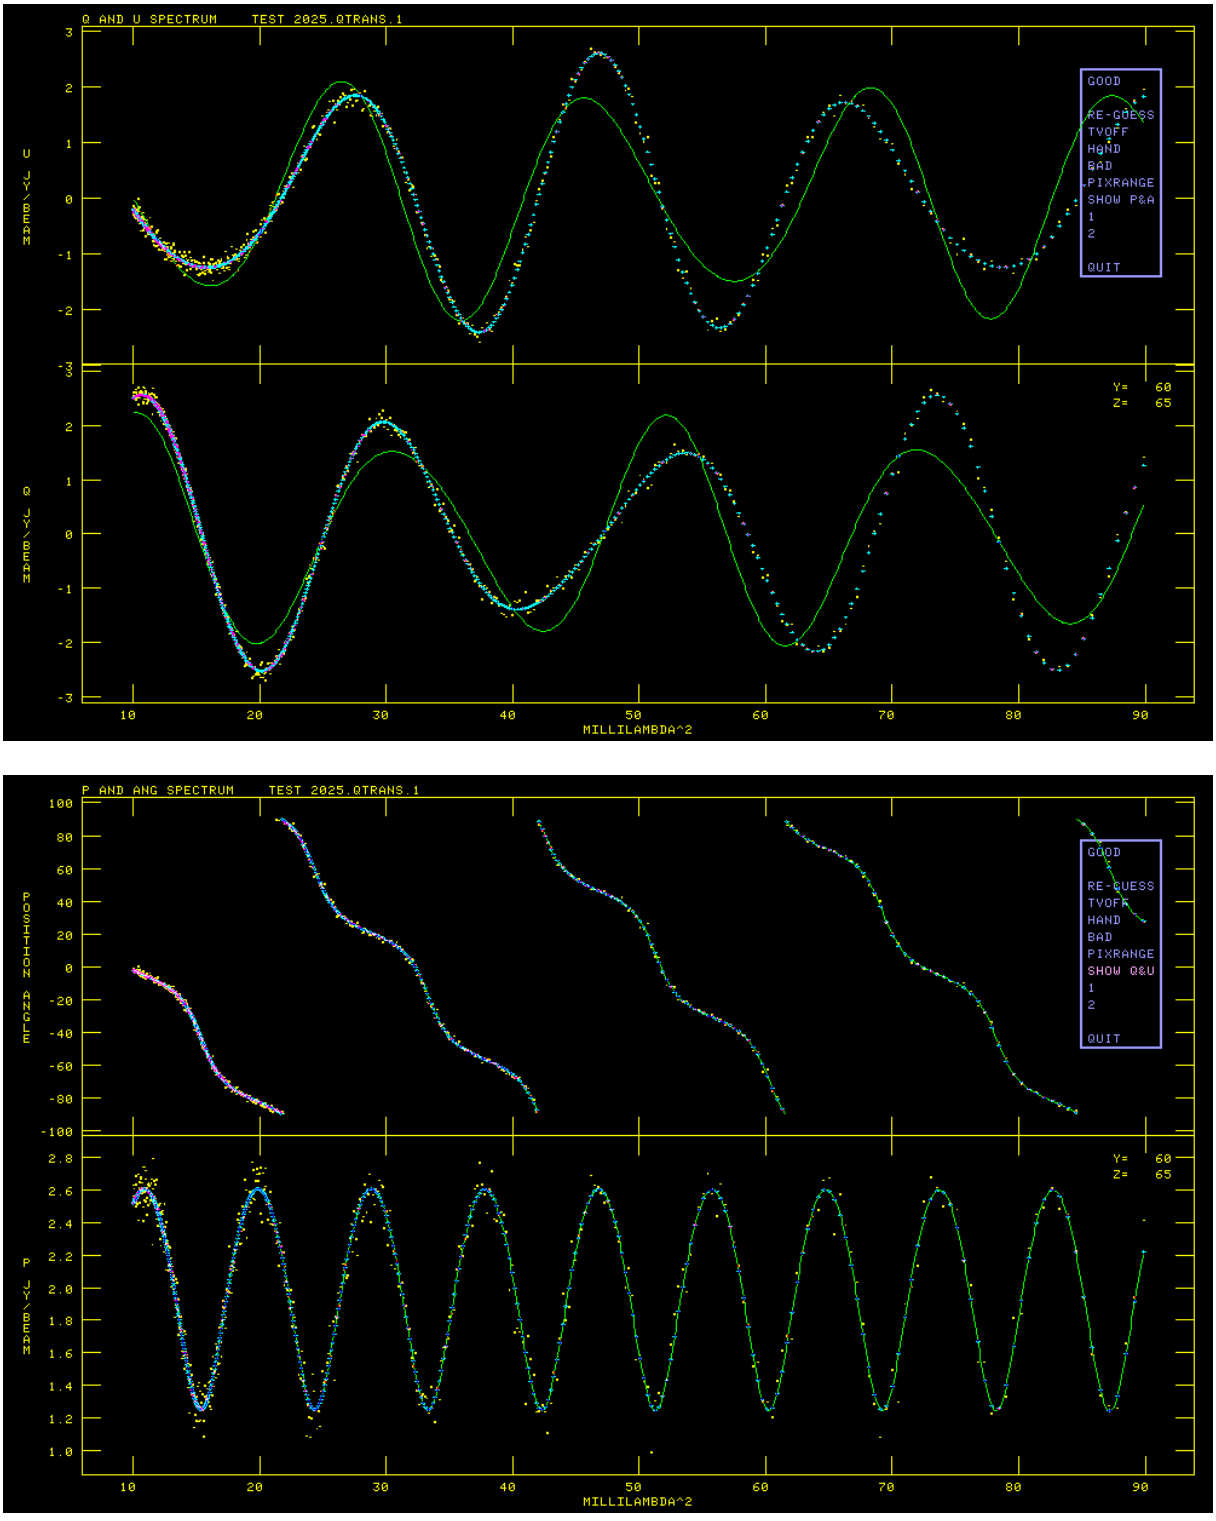

Figure 22: Top: Model field first pixel Q and U spectra, fit with two components. Initial guess not as close to fit, but fit now matches the data. Bottom: Spectrum re-plotted as amplitude and phase.

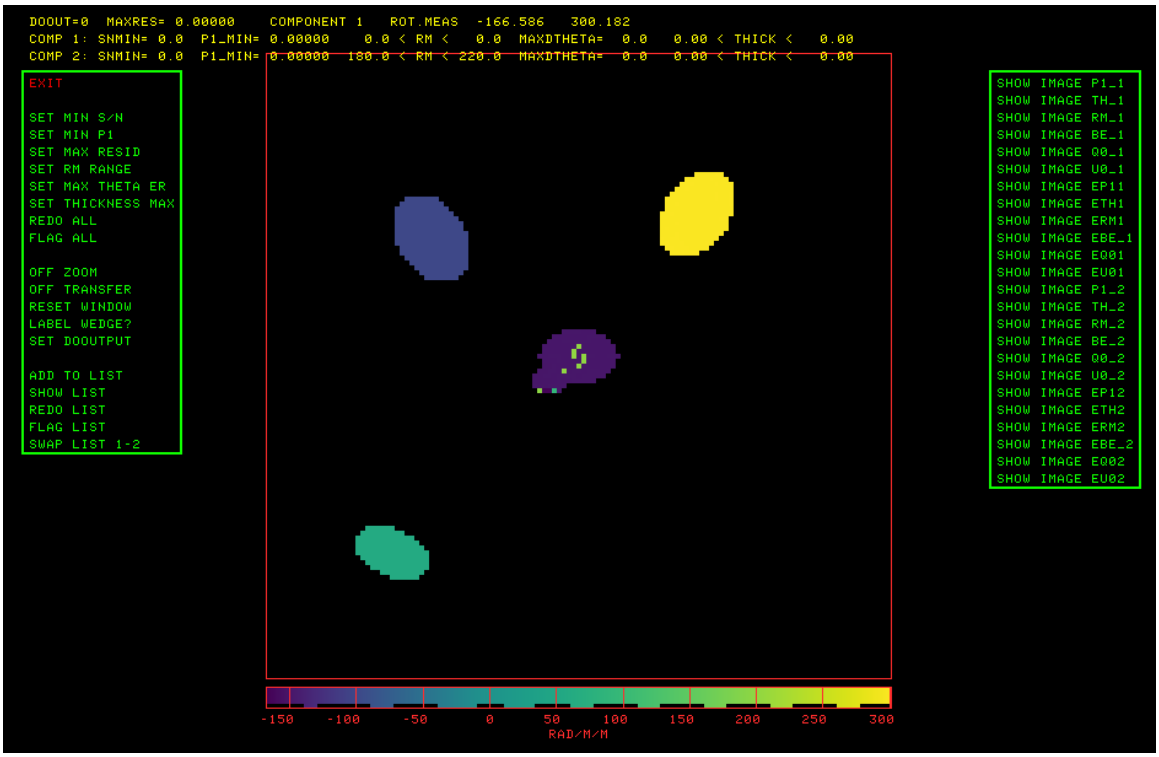

Figure 23: Model field showing rotation measure for component one. There are pixels needing correction. It has been colored by OFMCOLOR, the menu is labeled, and the acceptable range for component two has been set.

than RMSLIMIT. When that happens, you are shown the offending fit parameters and the task resumes with the plot at the bottom of Figure 21, already prompting you to enter a new guess, to allow you to try to fix things.

4.3 Editing and output

Eventually all positions selected by PCUT, ICUT, BLC, and TRC will have been fit. At this point, the task computes images of the fit parameters plus the Q and U polarization at zero wavelength and their uncertainties. It then offers a lengthy menu of options which will allow you to view these images and revisit positions that seem to have produced incorrect fits. If NGAUSS > 1, options to swap portions of image n with corresponding portions of image m are also offered. This "edit" menu is illustrated in Figure 23 for the NGAUSS = 2 and DOSPIX = 3 case. RMFIT will replicate pixels in both directions if needed to make small images large enough to see, but it can also handle images larger than the TV display area. The SET SPINX RANGE appears only when DOSPIX = 1 and the SET THICKNESS MAX appears only when DOSPIX = 2, 3, or 4, the first column of the menu shows the editing functions.

There are three kinds of editing implemented here. In the first, the user establishes the parameter extrema which should be viewed as acceptable. The extrema currently set are shown in the title lines. Then, RMFIT may be told to flag all solutions not meeting these criteria, or, more profitably, to revisit those positions to see if a better fit can be obtained. The other editing methods are similar, but act on a list of pixel positions. These may be entered by typing in values, by clicking on suspect pixel positions in the CURVALUE function described below or by using the BLOTCH LIST function also described below. The contents of the list may be viewed, the solutions at the positions may be flagged, or they may be revisited to attempt for a better solution, or the solutions at the listed positions may be swapped between components n and m. The menu

will offer only appropriate swaps between components, thus 1 and 2 in our NGAUSS=2 example. Swapping may be required if RMFIT gets confused as to which component you want to call number 1 and which number 2, as illustrated in Figure 23. After the flagging, revisiting, or swapping, the list is cleared.

EXIT	Exit RMFIT, writing output images if DOOUTPUT is now > 0.
SET MIN S/N	Set minimum amplitude S/N(s) for okay solutions
SET MIN P1	Set minimum polarization at 1 GHz for okay solutions
SET MAX RESID	Set maximum residual for okay solutions
SET RM RANGE	Set rotation measure value range(s) for okay solutions
SET MAX THETA ER	Set maximum error(s) in polarization angle for okay solutions
SET SPINX RANGE	Set spectral index range(s) for okay solutions
SET THICKNESS MAX	Set thickness maximum value(s) for okay solutions
REDO ALL	Re-do all solutions which are not okay
FLAG ALL	Mark bad all solutions which are not okay
OFF ZOOM	Turn off TV zoom
OFF TRANSFER	Turn off black & white and color TV enhancements
RESET WINDOW	Display full view of current image
LABEL WEDGE?	Turn on/off labeling of step wedge
SET DOOUTPUT	Increment DOOUTPUT in loop 0-3
PIXRANGE	To enter new plot ranges with the terminal
ADD TO LIST	Type in output pixel coordinates to add to list for Q and U and for total polarization and angle
SHOW LIST	Display pixel coordinates in list
REDO LIST	Re-do solutions for all pixels in list
FLAG LIST	Flag solutions for all pixels in list
SWAP LIST 1-2	Swap solutions for components 1 and 2 for all pixels in list

The second (and potentially third) menu columns contain NGAUSS sets of functions

SHOW IMAGE P1_1	Enter image interaction with total polarization at 1 GHz of component 1
SHOW IMAGE TH_1	Enter image interaction with polarization position angle of component 1
SHOW IMAGE RM_1	Enter image interaction with rotation measure of component 1
SHOW IMAGE BE_1	Enter image interaction with RM thickness (β) of component 1
SHOW IMAGE SP_1	Enter image interaction with spectral index of component 1
SHOW IMAGE Q0_1	Enter image interaction with Q at 0 wavelength of component 1
SHOW IMAGE U0_1	Enter image interaction with U at 0 wavelength of component 1
SHOW IMAGE EP1_1	Enter image interaction with uncertainty in total polarization at 1 GHz
SHOW IMAGE ETH_1	Enter image interaction with uncertainty in polarization position angle
SHOW IMAGE ERM_1	Enter image interaction with uncertainty in rotation measure of component 1
SHOW IMAGE UBE_1	Enter image interaction with uncertainty in RM thickness (β) of component 1
SHOW IMAGE ESP_1	Enter image interaction with uncertainty in spectral index of component 1
SHOW IMAGE EQ0_1	Enter image interaction with uncertainty in Q at 0 wavelength of component 1
SHOW IMAGE EU0_1	Enter image interaction with uncertainty in U at 0 wavelength of component 1

The SP options will appear only if DOSPIX = 1 was specified and the BE options will appear only if DOSPIX = 2, 3, or 4 was specified.

When you select one of these functions most of the following operations will appear in yet another menu. This menu is illustrated in Figure 24. Only one of the LOAD AS options will appear, with the next one in the sequence offered when the current one has been invoked. The SWAP n-m options will appear as needed when NGAUSS > 1. The NEXT WINDOW option appears when needed to display an image too large to fit on the display screen. These options mostly invoke familiar functions from *AIPS* to control the FUNCTYPE used in loading the image to the display, to enhance the image intensities, to color the enhanced image intensities, and to zoom the display.

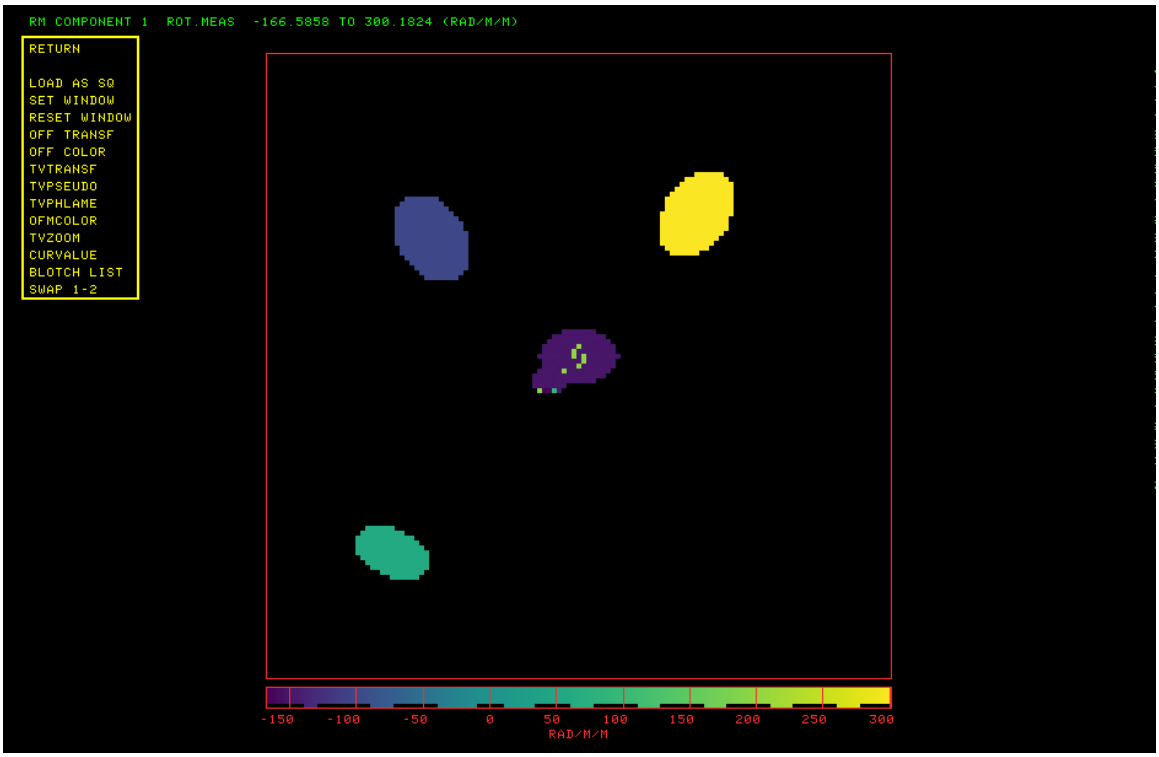

Figure 24: Image screen showing model field rotation measure of component 1. Pixels in the central "source" were solved incorrectly.

RETURN	Return to the above menus, image stays displayed
LOAD AS SQ	Re-load image with square root transfer function
LOAD AS LG	Re-load image with log transfer function
LOAD AS L2	Re-load image with extreme log transfer function
LOAD AS LN	Re-load image with linear transfer function
SET WINDOW	Set a sub-image to view
RESET WINDOW	Return to viewing the full image
OFF TRANSF	Turn off enhancement done with TVTRANSF
OFF COLOR	Turn off color enhancements
TVTRANSF	Black & white image enhancement
TVPSEUDO	Color enhancement of numerous sorts
TVPHLAME	Color enhancement of flame type, multiple colors
OFMCOLOR	Many standard color tables (OFMs) available
TVZOOM	Interactive zooming and centering of image
CURVALUE	Display value under cursor, mark pixels for list
BLOTCH LIST	Interactively add a blotch area to the pixel list
SWAP 1-2	Swap solutions for components 1 and 2 interactively
NEXT WINDOW	Move to next window in large images

Two operations in this menu are different from the usual. CURVALUE provides the capability of selecting positions for the edit "list." During the CURVALUE operation position the cursor over the desired pixel and press buttons A, B, or C to add that pixel to the list. The BLOTCH LIST and SWAP n-m operations use a TV blotch operation like that in the *AIPS* verb TVSTAT and task BLSUM. You are to mark with a "blotch" region those pixels in the present image which are to be added to the pixel list or to have their solutions swapped with those of the selected component. Instructions will appear in the message window as you proceed.

Begin by positioning the cursor at a pixel to be the first vertex of a connected sequence of vertices and press TV button A. Move to the next vertex and press button A again and repeat until you have marked all vertices for this region. Then press button D if you are done with this region or button C if you need to re-position one of the vertices. In this case, move the cursor to the vertex to be re-positioned, press button A and drag the vertex to the corrected position. Press button A or B to fix that vertex and go on to reset another vertex or D to end this region and list or swap the solutions. You may do this as many times as needed.

When you have finished getting the images just the way you want them, you may write them out as *AIPS* image files. Select the SET DOOUTPUT option until its value, shown at the top of the screen, is what you want. In RMFIT, values 1 and 3 cause residual image cubes to be written, while values 2 and 3 cause images of the parameter values and their uncertainties to be written. The Q and U residual images get classes QRESID and URESID, respectively. The polarization at 1 GHz, polarization angle, rotation measure, and spectral index or thickness (if any) images and their uncertainties get classes PPOLn, THETAn, ROTMEn, SPIXn, THICKn, DPPOLn, DTHETn, DROTMn, DSPIXn, and DTHICn, respectively, and Q and U at zero wavelength get classes QO_n, UO_n, DQO_n, and DUO_n, where n is the component number.

5 Copying the XG and ZE tables

The 31DEC25 versions of XGAUS and ZEMAN build the XG and ZE tables, respectively, with a fixed maximum number of components set when the tables are first created. These tables can be large files. A 32-component XG table will be the same size as a 200-channel spectral cube. Therefor, creating these tables with only the required number of components is probably a good idea. However, after having spent a good deal of effort in a good table, it would be wise the protect that effort by being able to copy the table either increasing or decreasing the maximum number of components.

XG2XG is a new task in 31DEC25 to do just this. It has only a few parameters including the file selection adverbs, INVERS to select the input table version, NGAUSS to select the number of components in the new table, and INEXT to specify a ZE table with any other value choosing an XG table. XG2XG always writes to a new table. If the output number of components is less than the input, XG2XG will check for, and report, any pixels for which an apparently good solution was lost.

6 Post-fit plotting

The images produced by XGAUS, ZEMAN and RMFIT may be displayed using all the usual tools such as KNTR, PROFL, and numerous other tasks. However, the display of the spectral data and the various fits to them required new tasks.

XG2PL plots a spectrum for a single pixel or for a rectangular or circular region about a single pixel. For each pixel included in the average, the task reads the I polarization image to obtain the data and the appropriate line(s) in the XG table to obtain the XGAUS solution for that pixel. It then computes the spectrum of each component in the model, plus the sum of the components, and the residual (data-model). Each of these then enter into the average of that parameter. Finally, the task plots a user-selected number of the parameters and, optionally, prints all of them to a text file. The plot may appear on the TV or be placed in a standard plot file attached to the I polarization image.

Optionally, XG2PL will also add the spectrum of the V polarization data and the results of the fitting done by ZEMAN at the same pixels as the I polarization Gaussians. The task reads the V polarization image for the data spectrum and the ZE table for the Zeeman-splitting model (either using the Gaussians or the simpler ones using the I polarization data). It computes the spectrum of each component (including the gain term in each), the net model (sum of the components but including the gain term only once), and the residual. Finally, the task plots a user-selected number of the parameters and, optionally, prints all of them to a text file. In general, the I polarization spectrum appears in the upper part of the plot and the V polarization

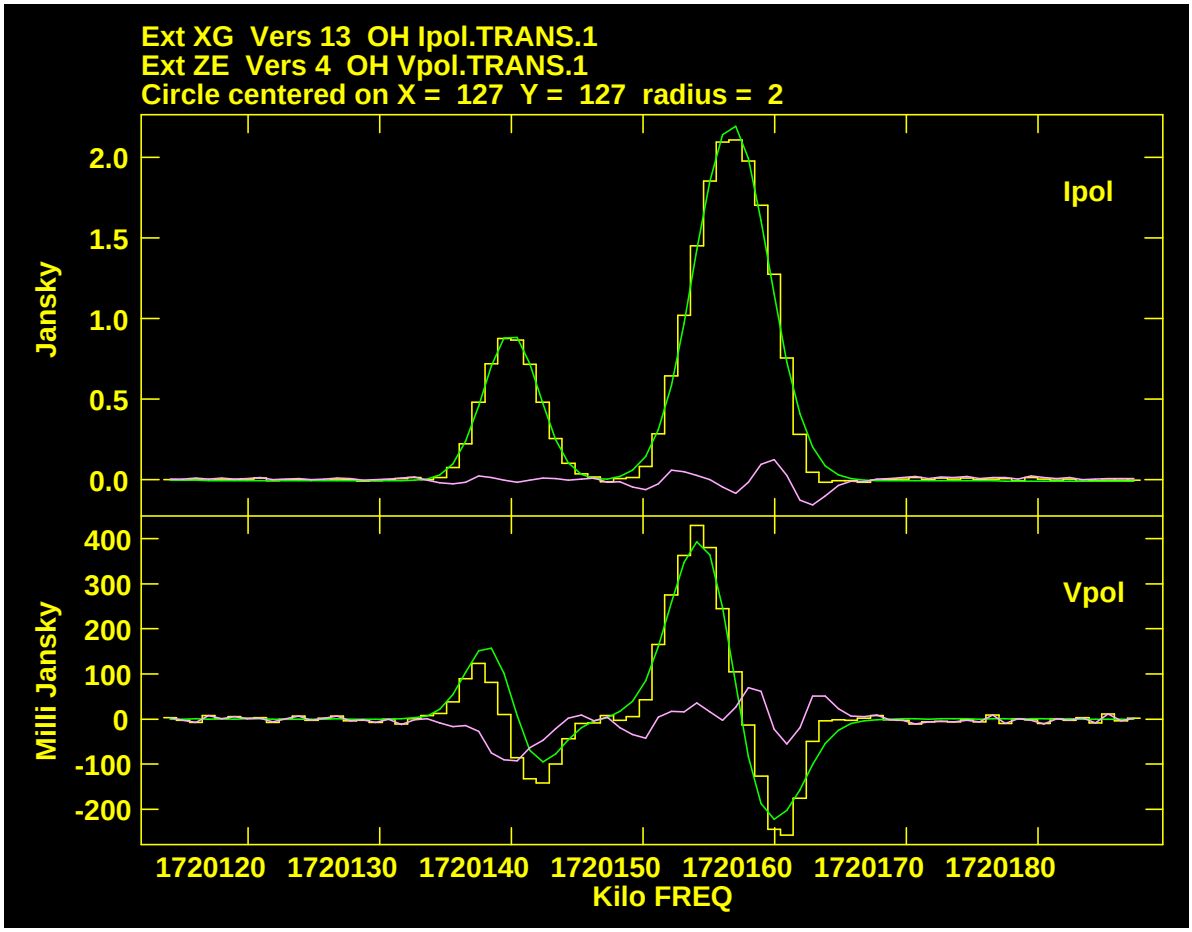

Figure 25: XG2PL example spectrum from a circular area for a Gaussian and Zeeman fit.

spectrum appears in the lower portion. Either portion may be omitted under control of the adverbs. The output of XG2PL is illustrated in Figure 25.

The inputs for XG2PL begin with the I polarization image which is required and then the V polarization image which may be omitted. INVERS and IN2VERS give the version numbers of the XG and ZE tables, respectively. APARM provides the central pixel coordinates, the plot intensity ranges for I and for V, the size of the rectangle (or circle) over which to average, and a flag limiting the average to those pixels having a model fit, BPARM is a set of flags selecting which parameters are plotted. CPARM selects the channel range to be plotted, the relative size of the V and I plots, the type of the horizontal axis (channels, frequency, velocity), whether the channels are plotted in reverse order, and whether the data are plotted with stepped or directly connected lines. OUTTEXT specifies the output text file, if any. The usual XYRATIO, LTYPE, DOTV, and GRCHAN adverbs control the scale, labeling, choice of TV versus plot file, and, if TV, which graphic channel(s) are used. If GRCHAN is zero, graphics channel 1 is used for data and labeling (usually yellow), channel 2 is used for the full model (usually green), channel 3 is used for the residual (usually pink), and channel 4 is used for the model component(s) (usually cyan).

RM2PL plots a spectrum for a single pixel or for a rectangular or circular region about a single pixel. For each pixel included in the average, the task reads the Q and U polarization images to obtain the data and the appropriate line(s) of the RM table to obtain the RMFIT solution for that pixel. It then computes the spectrum of each component in the model, plus the sum of the components, and the residual (data-model). Each of these then enter into the average of that parameter. Finally, the task plots a user-selected number of the parameters and, optionally, prints all of them to a text file. The plot may appear on the TV or be placed in

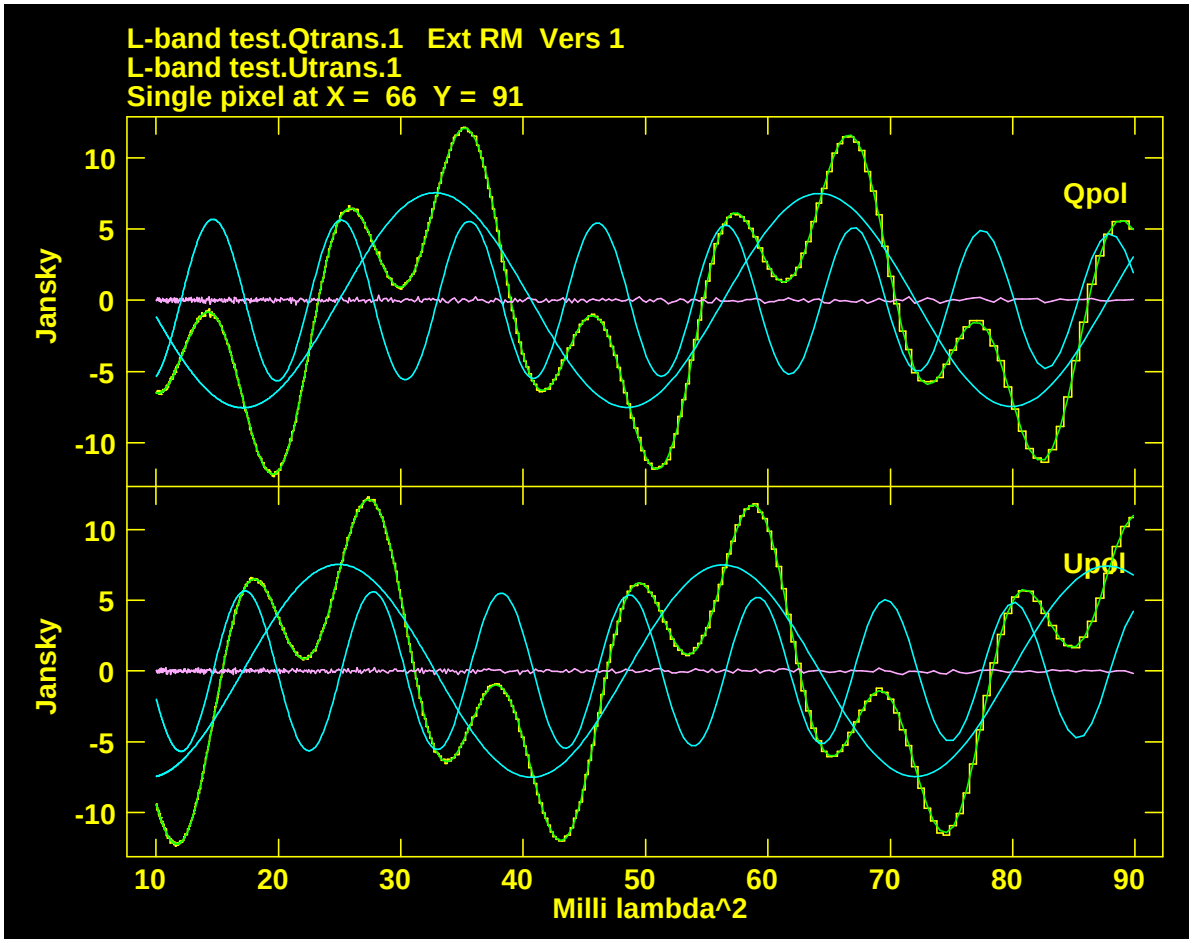

Figure 26: RM2PL example spectrum from a circular area for a 2-component rotation measure fit.

a standard plot file attached to the Q polarization image. The user may choose to omit all of the Q or all of the U polarization plot. The output of RM2PL is illustrated in Figure 26.

The inputs for RM2PL begin with the Q polarization image which is required and then the U polarization image which may be omitted. INVERS gives the version number of the RM table. APARM provides the central pixel coordinates, the plot intensity ranges for Q and for U, the size of the rectangle (or circle) over which to average, and a flag limiting the average to those pixels having a model fit. BPARM is a set of flags selecting which parameters are plotted. CPARM selects the channel range to be plotted, the relative size of the U and Q plots, the type of the horizontal axis (channels, frequency, wavelength, wavelength squared), whether the channels are plotted in reverse order, and whether the data are plotted with stepped or directly connected lines. OUTTEXT specifies the output text file, if any. The usual XYRATIO, LTYPE, DOTV, and GRCHAN adverbs control the scale, labeling, choice of TV versus plot file, and, if TV, which graphic channel(s) are used. If GRCHAN is zero, graphics channel 1 is used for data and labeling (usually yellow), channel 2 is used for the full model (usually green), channel 3 is used for the residual (usually pink), and channel 4 is used for the model component(s) (usually cyan).

7 Model creation

7.1 Visibility data models

There are two tasks which will modify an existing visibility ("uv") data set, adding models and noise to the existing data or replacing those data. UVMOD can apply a full set of calibration and data selection adverbs to the input uv data set. Then the data are scaled by FACTOR, which can be zero to eliminate the input data entirely. Finally, a model is added. The model can be a simple set of up to 4 NGAUSS components specified by adverbs CTYPE, FMAX, FPOS, and FWIDTH. Alternatively, adverb INLIST may specify a text file containing up to 9999 components. Each non-comment line in this file specifies

1.	I-polarization flux (Jy at header frequency)
2.	X-shift from reference (arc seconds)
3.	Y-shift from reference (arc seconds)
4.	Major axis (arc seconds)
5.	Minor axis (arc seconds)
6.	Position angle (degrees CCW from North)
7.	Type code: 1 Gaussian, 2 solid disk, 3 solid rectangle, 4 optically thin sphere, 5 exponential, else point
8.	Spectral index (see note below)
9.	Spectral index curvature
10.	Q-polarization flux (Jy at header frequency)
11.	U-polarization flux (Jy at header frequency)
12.	V-polarization flux (Jy at header frequency)

where trailing zeros may be omitted. If spectral index is used, note the definition of spectral index and flux. If x is the logarithm, base 10, of the frequency divided by the header frequency, then

$$F = F_0 10^{(\alpha x + \beta x^2)}$$

where α is the spectral index and β is its curvature. Thus the flux you must specify in the model is the flux at the frequency of the data.

SPMOD is a very similar task designed to do spectral-line modeling rather than spectral index and polarization. In this case INLIST is required and each of the up to 9999 non-comment lines contains

1.	RR-polarization flux (Jy)
2.	X-shift from reference (arc seconds)
3.	Y-shift from reference (arc seconds)
4.	Major axis (arc seconds)
5.	Minor axis (arc seconds)
6.	Position angle (degrees CCW from North)
7.	Type code: 1 Gaussian, 2 solid disk, 3 solid rectangle, 4 optically thin sphere, 5 exponential, else point
8.	Spectral line center (channels, but see note below)
9.	Spectral line width (FWHM channels)
10.	LL-polarization flux (Jy)

where trailing zeros may be omitted. The LL flux will be set equal to the RR flux if it is omitted. But note that both the RR and LL flux may be specified to have any value positive, zero, or negative. This will allow Zeeman-splitting to be modeled with any degree of splitting. Spectral channels are counted from BCHAN through ECHAN as 1 to N in the first spectral window (BIF), then $N+1$ through $2N$ in the next spectral window, and so forth. This allows you to put lines in any spectral window you desire.

7.2 Image data models

There are three tasks designed to modify an image adding a specified model and noise. `IMMOD` is intended for continuum images. The existing image may be scaled (or eliminated) and noise added. It adds `NGAUSS` components up to four using adverbs `OPCODE` to specify type, and `FMAX`, `FPOS`, and `FWIDTH` to specify peak brightness and pixel position and size. Alternatively, adverb `INLIST` may be used to specify up to 9999 components, one per line in the text file. Each non-comment line specifies

1.	Peak brightness (Jy/beam)
2.	Component X center (pixels)
3.	Component Y center (pixels)
4.	Major axis (pixels)
5.	Minor axis (pixels)
6.	Position angle (degrees CCW from North)

The more recent task `MODIM` is specifically designed for polarization modeling and has been used extensively in testing `RMFIT`. It takes in three image cubes, one for I polarization, one for Q, and one for U. Alternatively, it can create new images with no input images used. Up to 9999 components are taken from the text file specified by `INLIST`. Each non-comment line in this file specifies *every one of*

1.	I-polarization brightness (Jy/beam at 1 GHz)
2.	Q-polarization brightness (Jy/beam at 1 GHz)
3.	U-polarization brightness (Jy/beam at 1 GHz)
4.	Spectral index (see note below)
5.	Rotation measure (radians m^{-2})
6.	Rotation measure thickness (radians m^{-2})
7.	Component X center (pixels)
8.	Component Y center (pixels)
9.	Major axis (pixels)
10.	Minor axis (pixels)
11.	Position angle (degrees CCW from North)
12.	Type code: 1 Point, 2 Gaussian, 3 solid disk, 4 solid rectangle, 5 optically thin sphere, 6 exponential, else point

Note that all parameters must be specified for every component. Adverb `OPTYPE` specifies the type of rotation measure thickness model, with values `'SLAB'`, `'GAUS'`, and `'EXP '`. The brightness of a component at frequency ν in GHz is $F = F_1\nu^{\alpha}$. Input (if any) and output images are in the transposition that places celestial coordinates first and frequency on the third axis.

A third modeling task, `MODSP`, models I and V image cubes from models similar to those of `MODIM`, but suitable for spectral lines. It takes in two image cubes, one for I or RR polarization and one for V or LL polarization. Alternatively, it can create new images with no input images used. Up to 9999 components are taken from the text file specified by `INLIST`. Each non-comment line in this file specifies

1.	R-polarization brightness at line center (Jy/beam)
2.	L-polarization brightness at line center (Jy/beam)
3.	Component X center (pixels)
4.	Component Y center (pixels)
5.	Major axis (pixels)
6.	Minor axis (pixels)
7.	Position angle (degrees CCW from North)
8.	Type code: 1 Point, 2 Gaussian, 3 solid disk, 4 solid rectangle, 5 optically thin sphere, 6 exponential, else Gaussian
9.	Line center (channels)
10.	Line width (FWHM of Gaussian in channels)
11.	Change of line center per X pixel (channels/pixel)
12.	Change of line center per Y pixel (channels/pixel)
13.	Change of line width per X pixel (channels/pixel)
14.	Change of line width per Y pixel (channels/pixel)

Note that parameters 1–10 must be specified for every component. The remaining position and width derivatives will be taken as zero if omitted. Input (if any) and output images are in the transposition that places celestial coordinates first and frequency on the third axis. Input images may be in either I and V polarizations or RR and LL, output images will be in I and V. Input images are scaled by `FACTOR` before the model is added. `MODSP` allows the spectral channel parameters to be linear functions of position around the center of the component. Thus, the line center channel and width are given by

$$\begin{aligned} C(x,y) &= C_i + \frac{dC_i}{dx}(x - X_i) + \frac{dC_i}{dy}(y - Y_i) \\ W(x,y) &= W_i + \frac{dW_i}{dx}(x - X_i) + \frac{dW_i}{dy}(y - Y_i) \end{aligned}$$

where, for component i, C_i is parameter 9 above, $\frac{dC_i}{dx}$ is parameter 11, $\frac{dC_i}{dy}$ is parameter 12, X_i is parameter 2, Y_i is parameter 3. W_i is parameter 10, $\frac{dW_i}{dx}$ is parameter 13, and $\frac{dW_i}{dy}$ is parameter 14.