

Modeling Spectral Cubes in *AIPS*

Eric W. Greisen

May 6, 2020

Abstract

AIPS has done Gaussian fitting along the x -axis of image cubes with task *XGAUS* since the 1980s. That task has recently been overhauled to be much easier to use and much more capable. In like fashion, new tasks *ZEMAN* and *RMFIT* have been 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 (*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*.](#)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 *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.

¹Color is used to indicate changes made to this document: the June 2014 edition, [the September 2015 revision](#), [the November 2017 revision](#) and [the Ma 2020 revision](#).

²Brentjens, M. A. and de Bruyn, A. G. 2005, *Astron. & Ap.*, 1217.

The author undertook a re-write of XGAUS 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. Two new 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.

Two new tasks have been written to prepare plot files from the fits done by these tasks. 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

$$T(x, y, z) = A + Bx + \sum_1^{\text{NGAUSS}} T_i e^{-\alpha(x-x_i)^2/\sigma_i^2} \quad (1)$$

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 with 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 so leave it zero at this point. Set DOTV = 2 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, but, if you are fitting weak Gaussians in the presence of very strong ones, you might wish to cut off the highest values.

(This may make it harder to set initial guesses however since you can't point at the peak of the 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. It may be changed if you re-start on a pre-existing XG table since those tables contain room for the maximum number of Gaussians allowed (8 at this writing). 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. 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, either in your AIPS terminal window (D0TV = 1) or, as shown in the present figures, on the TV (D0TV = 2). The first spectrum to be fit (from an OH 1720 MHz maser source) is illustrated in 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	
RE-GUESS	E
BAD	B
QUIT	Q

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 (or Q on the terminal) causes the task to quit at this point. You may re-start later. The option BAD (B on the terminal) 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 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




Figure 1: First OH spectrum to fit, initial “guess” is all zero.

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 3 (usually pink). This is illustrated in Figure 3. A different menu appears at this point containing

GOOD	
DO FIT	D
RE-GUESS	E or R
TVOFF	T
HAND	H
BAD	B
1	1
2	2
QUIT	Q

The option to QUIT (or Q on the terminal) causes the task to quit at this point. You may re-start later. The option BAD (B on the terminal) will mark this position as failed and go on to the next position. The option RE-GUESS (E or R on the terminal) will loop back to prompt you for a new guess and repeat the fit. Options 1, 2, . . . , NGAUSS will loop back to plot an initial guess with the selected number of Gaussians. (NGAUSS=2 in the




Figure 2: First OH spectrum to fit: better initial “guess” has been entered.

current example.) Option HAND (H on the terminal) 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 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 with the plot of Figure 3 to allow you to try to fix things. Option GOOD (any other initial character on the terminal) 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. It then offers




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

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 OH 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. Note too that, for legibility in all figures in this memo, the *ALPS* TV was run with double-sized characters.

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, 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 or by clicking on suspect pixel positions in the CURVALUE function 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. After the flagging, revisiting, or swapping, the list is cleared. The first column of the menu shows the following:




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

EXIT	Exit XGAUS, writing output images if D00OUTPUT 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
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 D00OUTPUT	Increment D00OUTPUT 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 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
SWAP LIST 1-2	Swap solutions for components 1 and 2 for all pixels in list

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




Figure 5: OH maser field, edit view

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

On very crowded menus, the word `SHOW` may be omitted. When you select one of these functions most of the following operations will appear in yet another menu. This menu is illustrated in Figure 6. 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 impose your selection of component number n and component m in case `XGAUS` got confused. 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 6: HI galaxy: peak value image interaction

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
TVZOOM	Interactive zooming and centering of image
CURVALUE	Display value under cursor, mark pixels for 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 SWAP $n-m$ operation uses 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 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




Figure 7: HI galaxy: center channel image interaction

or B to fix that vertex and go on to reset another vertex or D to end this region and 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 *ATPS* image files. Select the SET D00OUTPUT option until its value, shown at the top of the screen, is what you want. In XGAUS, values 1 and 3 cause a residual image cube to be written, while values 2 and 3 cause images of the parameter values and their uncertainties 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 AMPL n , CENTR n , WIDTH n , DAMPL n , DCENT n , and DWIDT n , and the flux and its uncertainty get class FLUX n , and DFLUX n , where n is the component number.

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

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

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.³ The *ATPS* 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

³Sault, R.J., Killeen, N. E. B., Zmuidzinas, J., Loushin, R. 1990, *Ap. J.*, **74**, 437-461.

applied, another function may be used instead:

$$V(x, y, z) = A(y, z)I(x, y, z) + 0.5 \sum_1^{\text{NGAUSS}} B_i(y, z) \frac{dG_i(x, y, z)}{dx} \quad (3)$$

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 the example shown below, the present data lack the spatial resolution to separate two spectral components, but this second method easily finds two different magnetic field values. This source has been observed with much better spatial resolution which separates the two components. The present field values match those found in the published, high-resolution image. 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 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

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

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. D00UTPUT controls what files are written — this may be changed interactively so leave it zero at this point. Set D0TV = 2 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.

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, either in your AIPS terminal window (D0TV = 1) or, as shown in the present figures, on the TV (D0TV = 2). The first spectrum to be fit (from an OH 1720 MHz maser source) is illustrated in Figure 8 showing that the first guess has zero for the two $B_i(y, z)$ but a reasonable guess at the $A(y, z)$. Thus, the green line mimics the shape of the I spectrum, not the V spectrum. The menu that appears at this point is




Figure 8: First OH spectrum to fit, initial “guess” has Gaussians but the two $B_i(y, z)$ are zero.

DO FIT	
BAD	B
QUIT	Q

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 (or Q on the terminal) causes the task to quit at this point. You may re-start later. The option BAD (B on the terminal) 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 9. The second pixel starts with the $B_i(y, z)$ of the last fit (Figure 10) and so is a much better initial guess.




Figure 9: First OH spectrum to fit, plot after fitting.




Figure 10: Second OH spectrum to fit, initial “guess” has Gaussians and values for $B_i(y, z)$.




Figure 11: ZEMAN fit (component 1 splitting) and image editing screen.

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 fit parameters $A(y, z)$ and $B_i(y, z)$ and their uncertainties. 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 in Figure 11. Note that the size of the OH 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.

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. 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. The other editing method is similar, but acts on a list of pixel positions. These may be entered by typing in values or by clicking on suspect pixel positions in the CURVALUE function 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. The first column of the menu includes:

EXIT	Exit XGAUS, writing output images if DDOUTPUT 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 DDOUTPUT	Increment DDOUTPUT 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 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

The second menu column 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 F_n and EF_n 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 12, these options are

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
TVZOOM	Interactive zooming and centering of image
CURVALUE	Display value under cursor, mark pixels for list
NEXT WINDOW	Move to next window into 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. 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.

Figure 12: ZEMAN image of $B_1(y, z)$ with display options.

When you have finished getting the images just the way you want them, you may write them out as *ALPS* 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, and 7 cause a residual (data-model) image cube to be written, while values 2, 3, 6, and 7 cause images of the gain and field values and their uncertainties to be written and values 1, 3, 5, and 7 cause an image of the V cube to be written with the gain times the I-polarization cube subtracted. 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 FIELD n and DFELD n , with $n = 1$ to the maximum NGAUSS.

4 Polarization fitting

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:




Figure 13: 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 .

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

“Faraday Rotation Measure Synthesis” is then the Fourier analysis of Equation 5 as initially described by Brentjens and de Bruyn.⁴ 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 13. 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

⁴Brentjens, M. A. and de Bruyn, A. G. 2005, *Astron. & Ap.*, 1217.

is one of

$$\begin{aligned}
 F(\beta_j, \lambda(i)^2) &= \left(\frac{\lambda_1^2}{\lambda(i)^2} \right)^{\frac{\beta_j}{2}} \\
 &= \frac{\sin(\beta_j \lambda(i)^2)}{\beta_j \lambda(i)^2} \\
 &= e^{-\ln(2) \left(\frac{\beta_j \lambda(i)^2}{1.8954} \right)^2} \\
 &= e^{-\ln(2) \frac{\beta_j \lambda(i)^2}{1.8954}}
 \end{aligned}$$

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. Set DOTV = 2 to use TV menus to prompt you. The badly named adverb NGAUSS controls the maximum number of components to fit and is limited to 4. 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. Then it reads the table every YINC rows and ZINC planes and, for those with brightnesses exceeding PCUT and ICUT,




Figure 14: Model field first pixel FARS rotation measure spectrum to fit.

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, either in your AIPS terminal window (DOTV = 1) or, as shown in the present figures, on the TV (DOTV = 2). The first spectrum to be fit is illustrated in Figure 14 showing that the first guess for more than one component is not useful. The menu that appears at this point is

DO FIT	D
RE-GUESS	E
BAD	B
QUIT	Q

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 (or Q on the terminal) causes the task to quit at this point. You may re-start later. The option BAD (B on the terminal) 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 14 is RE-GUESS which causes the task to prompt you first to “Position cursor at center (RM) of component 1”. Move the cursor to the peak of component 1 and press any TV button. This selects the rotation measure value of component 1. This prompt is repeated for components two through NGAUSS. To omit the fit for a particular component at this position, move the cursor off the plot (but still inside the TV window) before pressing the TV button for that component.

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 two components initially, getting the unhappy result shown in Figure 15.

GOOD	
DO FIT	D
RE-GUESS	E or R
TVOFF	T
HAND	H
BAD	B
PIXRANGE	P
SHOW Q&U	S
SHOW P&A	S
1	1
2	2
3	3
QUIT	Q

The menu that appears at this point offers the option to QUIT (or Q on the terminal) which causes the task to quit at this point. You may re-start later. The option BAD (B on the terminal) will mark this position as failed and go on to the next position. The option RE-GUESS (E or R on the terminal) will loop back to prompt you for a new guess and repeat the fit. Options 1, 2, . . . , NGAUSS will loop back to plot an initial guess with the selected number of Gaussians. Option HAND (H on the terminal) 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 15 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 GOOD (any other initial character on the terminal) tells the task that you are (reasonably) happy and that it should go on to the next position.

In the present case, the user has wisely decided to enter a new guess, this time specifying all three components. The happy result is shown in Figure 16. *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 17.* Option




Figure 15: Model field first pixel Q and U spectra, fit with only two components (there are actually three).

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 with the plot of Figure 18, already prompting you to enter a new guess, to allow you to try to fix things.




Figure 16: Model field first pixel Q and U spectra, fit with three components.

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 $\text{NGAUSS} > 1$, options to swap portions of image n with corresponding portions of image m are also offered. This “edit” menu is illustrated in Figure 19 for the $\text{NGAUSS} = 2$ and $\text{DOSPIX} = 3$ 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 $\text{DOSPIX} = 1$ and the SET THICKNESS MAX appears only when $\text{DOSPIX} = 2, 3$, 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 or by clicking on suspect pixel positions in the CURVALUE function 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 $\text{NGAUSS}=2$ 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 19. After the flagging, revisiting, or swapping, the list is cleared.




Figure 17: Model field first pixel total polarization and angle spectra, fit with three components.




Figure 18: Model field FARS rotation measure spectrum at coordinate for which the fit fails.




Figure 19: Model field, showing rotation measure for component **two**, for which there are pixels needing correction.

EXIT	Exit XGAUS, writing output images if DOOOUTPUT 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 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 DOOOUTPUT	Increment DOOOUTPUT 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 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
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




Figure 20: Model field rotation measure of component 1 showing pixels in the upper **and central** “sources” which were solved incorrectly.

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 of component 1
SHOW IMAGE ETH_1	Enter image interaction with uncertainty in polarization position angle of component 1
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 20. 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.

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
TVZOOM	Interactive zooming and centering of image
CURVALUE	Display value under cursor, mark pixels for list
SWAP 1-2	Swap solutions for components 1 and 2 interactively
NEXT WINDOW	Move to next window into 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 SWAP $n-m$ operation uses 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 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 swap the solutions. You may do this as many times as needed.

The images displayed were edited, first by selecting pixels in the image displayed in Figure 20 with CURVALUE and then re-fitting the spectra with REDO LIST. Additional editing was needed for poor solutions at the upper right of the central “source.” SET RM RANGE and REDO ALL were used. The SWAP 2-3 option was used to correct the lower left component. A more crowded edit menu is then shown for three components with DOSPIX = 3 in Figure 21.

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 DDOOUTPUT 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 PPOL n , THETAn, ROTMEN, SPIX n , THICK n , DPPOL n , DTHET n , DROTM n , DSPIX n , and DTHIC n , respectively, and Q and U at zero wavelength get classes Q0 $_n$, U0 $_n$, DQ0 $_n$, and DU0 $_n$, where n is the component number.

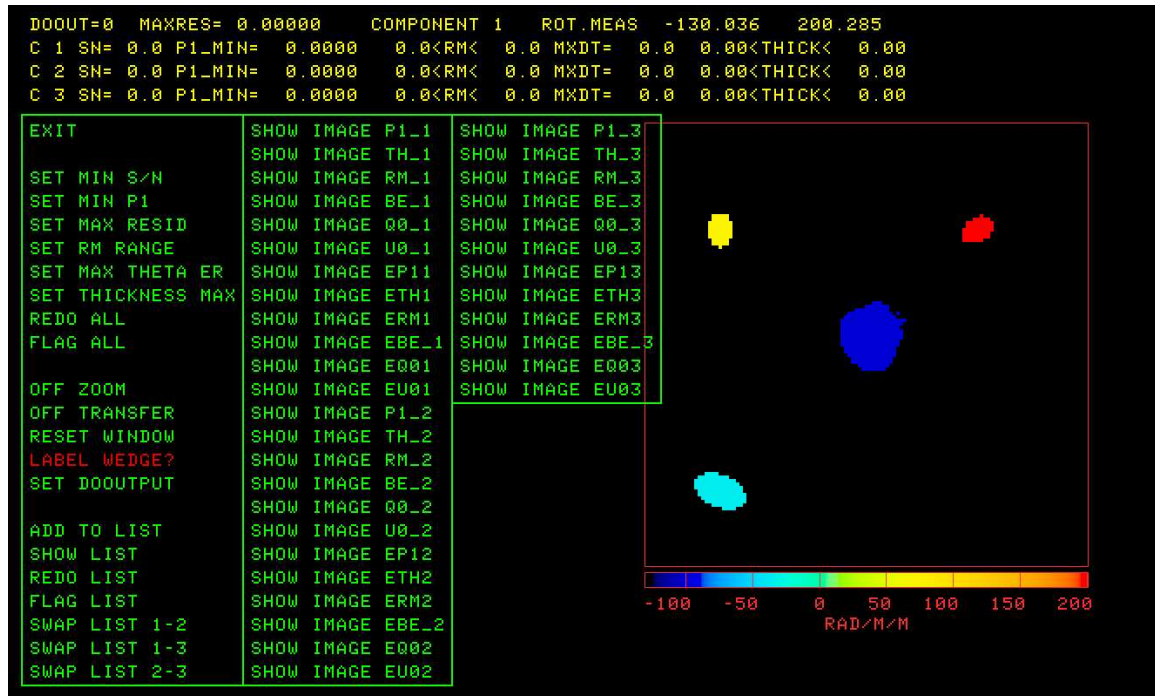


Figure 21: Model rotation measure of component 1 after editing showing edit menu for NGAUS = 3 and DOSPIX = 3.

5 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 on 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 spectrum appears in the lower portion. Either portion may be omitted under control of the adverbs. The output of XG2PL is illustrated in Figure 22.

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




Figure 22: Example spectrum from a circular area for a Gaussian and Zeeman fit.

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, DDTV, 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 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 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 23.

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




Figure 23: Example spectrum from a circular area for a 2-component rotation measure fit.

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).

6 Model creation

6.1 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.

6.2 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, similar to MODSP, has recently appeared. This task 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.