

Modeling Absorption-line Cubes in *AIPS*

Eric W. Greisen

May 6, 2020

Abstract

AIPS does Gaussian fitting of spectral lines with recently-overhauled task *XGAUS* and can fit V polarization image cubes for Zeeman-splitting with the relatively new task *ZEMAN*. Both of these tasks are designed for emission spectra in which the noise is not a function of spectral channel. In absorption, however, the noise in optical depth becomes high when the optical depth is high. Therefore, new tasks *AGAUS* and *ZAMAN* have been written to provide similar functions but with mathematics suitable for absorption lines. This memo describes the new tasks in some detail and includes a description of a new, simplified modeling task *MODAB* which may also be useful. That task has shown that the results of these four tasks are biased by the presence of the Zeeman splitting and need modest correction if they are meant to describe the actual pre-splitting line widths and magnetic field.¹

1 Introduction

XGAUS was 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 sets 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. *XGAUS* was extensively re-written beginning in 2013 and is now a quite friendly and powerful task. Accompanying this task is *ZEMAN* which first appeared in 2013 and which solves for Zeeman splitting of the right and left polarizations using a new Gaussian model or the older methods involving the derivative of the unpolarized spectra. These tasks are described in detail in *AIPS* Memo 118 (revised)²

However, these tasks are designed for emission lines for which the noise is not a significant function of spectral channel. In absorption, the optical depth is the physical parameter which one might expect to be Gaussian. However, the noise is a significant function of spectral channel whenever the optical depths become large. This calls for a new fitting routine to use Gaussian models in optical depth, but to fit the actual spectra for which the noise is approximately independent of spectral channel. This new task is named *AGAUS* and is very similar to *XGAUS*, but with enough differences that combining the two functions in one task is impractical.

A revision of *ZEMAN*, called *ZAMAN*, has also been written to solve for Zeeman splitting using the appropriate mathematics for absorption-line Gaussians. That mathematics will be presented below in some detail. A simple modeling task was written to prepare image cubes of known properties in order to test these new tasks. It has made it clear that the answers from *XGAUS* and *AGAUS* are affected by the Zeeman splitting as are the answers from *ZEMAN* and *ZAMAN*. The nature of the bias in the answers will be discussed below and correction methods suggested.

¹Color is used to indicate changes made to this document: the March 2017 edition and the May 2020 revision.

²Greisen, Eric W., 2013+, "Modeling Spectral Cubes in *AIPS*", *AIPS* Memo 118, NRAO, Socorro, NM.

2 Absorption-line Gaussian fitting: AGAUS

In absorption, we assume that the optical depth is proportional to the number of atoms/molecules in the line of sight and that the radial velocities of these absorbers is approximately Gaussian. Thus the function we seek to fit is

$$T(x, y, z) = (A(y, z) + B(y, z) * x) \times \exp(-\tau(x, y, z)) \quad (1)$$

where

$$\tau(x, y, z) = \sum_1^{\text{NGAUSS}} \tau_i e^{-\alpha(x-x_i)^2/\sigma_i^2} \quad (2)$$

where x is the coordinate along the first axis of the image cube, $A(y, z) + B(y, z)x$ is the spectral baseline (the continuum), $\tau_i(y, z)$ are the peak optical depths 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 .

AGAUS will fit this formulation to every row of a Stokes I image cube. Note that this cube must contain the full continuum emission of the source ($A(y, z)$) in order to be described by these equations. 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.

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 AGAUS. 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 taken as the peak value with no slope. ORDER controls the order of the baseline with ≤ 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. However, the plots are self-scaled from the data only and so may not include the residuals which should be near zero. Set PIXRANGE if you must see the residuals. In general it is best to set PIXRANGE to zero to see the full range of image values in best detail, but, if you are fitting weak Gaussians in the presence of very strong ones, you might wish to cut off the lowest values. This will not interfere with setting initial guesses as the plot of apparent optical depth is always self-scaled. 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 AGAUS or from your knowledge of the noise in your data cube.

Just as in XGAUS, AGAUS 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)..




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

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 (DOTV = 1) or, as shown in the present figures, on the TV (DOTV = 2). The first spectrum to be fit (from an OH absorption 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




Figure 2: First OH spectrum to fit: optical depth spectrum to mark the Gaussians.

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 change the plot to a spectrum of apparent optical depth, such as that shown in Figure 2. It 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) 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




Figure 3: First OH spectrum to fit: better initial “guess” has been entered.

in Figure 3. 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 4. The individual Gaussians are also plotted in a graphics plane which is only visible where no other graphics plane has been turned on.

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
3	3
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=3 in the




Figure 4: First OH spectrum to fit: good fit obtained.

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 optical depth), 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. AGAUS will then repeat the display in Figure 4 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 4 to allow you to try to fix things. Option GOOD (and 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, AGAUS 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.




Figure 5: OH absorption source peak of component 1 displayed on edit view

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 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, AGAUS 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, AGAUS 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, 2, and 3 in our $NGAUSS=3$ example. Swapping may be required if AGAUS 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:

EXIT	Exit AGAUS, writing output images if DDOUTPUT 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 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
SWAP LIST 1-2	Swap solutions for components 1 and 2 for all pixels in list
SWAP LIST 1-3	Swap solutions for components 1 and 3 for all pixels in list
SWAP LIST 2-3	Swap solutions for components 2 and 3 for all pixels in list

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

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 image interaction 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 AGAUS 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.

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

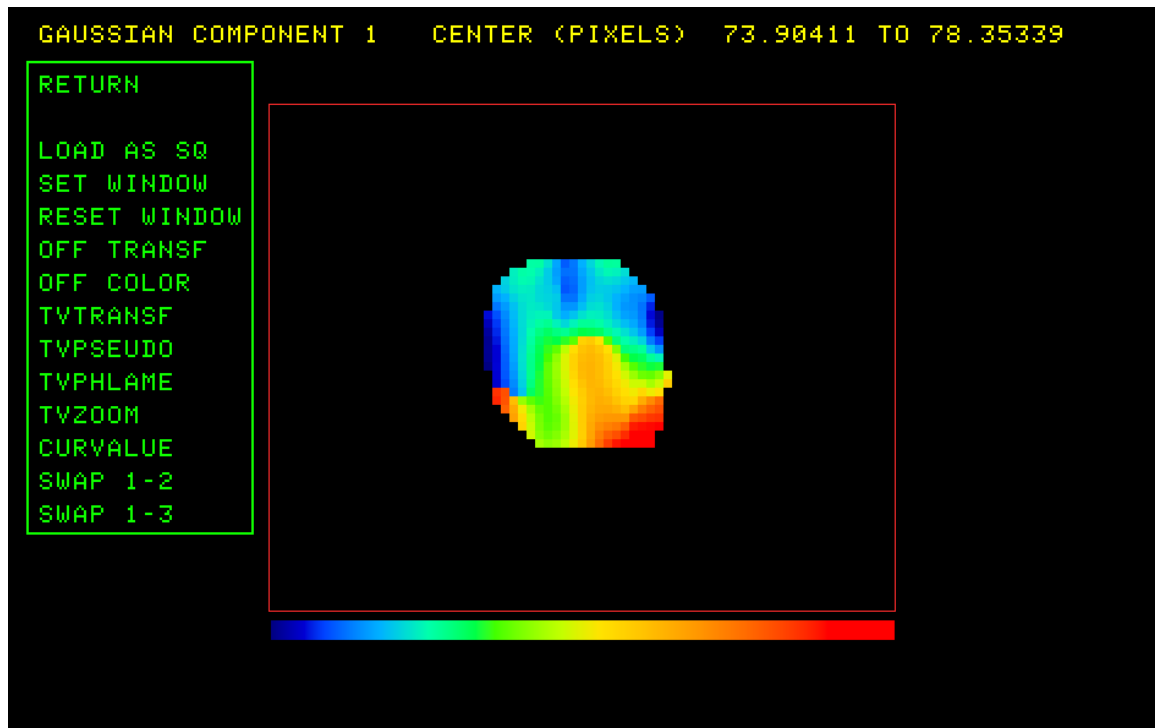


Figure 6: Center channel component 1 in image interaction stage

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.

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
SWAP 1-3	Swap solutions for components 1 and 3 interactively
SWAP 2-3	Swap solutions for components 2 and 3 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 AGAUS, 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: ZAMAN

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

$$V(x, y, z) = A(y, z)I(x, y, z) + 0.5B(y, z) \frac{dI(x, y, z)}{dx} \quad (3)$$

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 *AIPS* task ZAMAN 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 AGAUS applied, another function may be used instead. I present below a derivation of the formula for absorption, neglecting the leakage term until the end. Let us assume that the optical depth is due to a sum of Gaussians each having its own depth, center channel, and width. Thus

$$I(x, y, z) = C(y, z) \exp\left(-\sum_1^{\text{NGAUSS}} \tau_i(x, y, z)\right) \quad (4)$$

$$= C(y, z) \prod_1^{\text{NGAUSS}} \exp(-\tau_i(x, y, z)) \quad (5)$$

When shifted by magnetic fields, assumed different for each Gaussian, the formula for the V Stokes becomes

$$V(x, y, z) \equiv C(y, z) \left(\prod_1^{\text{NGAUSS}} \exp(-\tau_{iR}(x, y, z)) - \prod_1^{\text{NGAUSS}} \exp(-\tau_{iL}(x, y, z)) \right) \quad (6)$$

$$= C(y, z) \prod_1^{\text{NGAUSS}} \exp(-\tau_{iA}) \exp(-(\tau_{iR} - \tau_{iA})) - \prod_1^{\text{NGAUSS}} \exp(-\tau_{iA}) \exp(-(\tau_{iL} - \tau_{iA})) \quad (7)$$

where we assume that the difference between τ_{iL} and τ_{iA} is small. Then

$$\exp(-(\tau_{iL} - \tau_{iA})) \approx 1 + \tau_{iA} - \tau_{iL} \quad (8)$$

$$= 1 + \tau_{i0} \left(\exp(-\alpha(x - x_i)^2 / \sigma_i^2) - \exp(-\alpha(x - (x_i + \delta_i/2))^2 / \sigma_i^2) \right) \quad (9)$$

$$\approx 1 + \tau_{i0} \exp(-\alpha(x - x_i)^2 / \sigma_i^2) \left[1 - \exp(\alpha(x - x_i)\delta_i / \sigma_i^2 - \alpha\delta_i^2 / (4\sigma_i^2)) \right] \quad (10)$$

$$\approx 1 + \tau_{i0} \exp(-\alpha(x - x_i)^2 / \sigma_i^2) \left[-\alpha(x - x_i)\delta_i / \sigma_i^2 \right] \quad (11)$$

$$\exp(-(\tau_{iR} - \tau_{iA})) \approx 1 + \tau_{i0} \exp(-\alpha(x - x_i)^2 / \sigma_i^2) \left[\alpha(x - x_i)\delta_i / \sigma_i^2 \right] \quad (12)$$

Thus each of the terms in Equation 7 is of the form

$$= C(y, z) \prod_1^{\text{NGAUSS}} \exp(-\tau_{iA})(1 \pm a_i) \quad (13)$$

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

$$= C(y, z) \exp\left(-\sum \tau_i\right) \prod_1^{\text{NGAUSS}} (1 \pm a_i) \quad (14)$$

$$\approx C(y, z) \exp(-\tau) \left(1 \pm \sum a_i\right) \quad (15)$$

We note that the product of the $\exp(-\tau_{iA})$ is just the total absorption factor $\exp(-\tau)$ and the product of a number of $(1 \pm a_i)$, where a_i is small, is approximately $1 \pm \sum a_i$. We end up with

$$\tau(x, y, z) = \sum_1^{\text{NGAUSS}} \tau_{i0} \exp(-\alpha(x - x_i)^2 / \sigma_i^2) \quad (16)$$

$$V(x, y, z) \approx C(y, z) \exp(-\tau(x, y, z)) \sum_1^{\text{NGAUSS}} (2\tau_i \alpha (x - x_i) \delta_i / \sigma_i^2) \exp(-\tau_i(x, y, z)) \quad (17)$$

$$V(x, y, z) = A(y, z) I(x, y, z) + 0.5 \sum_1^{\text{NGAUSS}} B_i(y, z) \frac{-d\tau_i(x, y, z)}{dx} \quad (18)$$

where $\tau_i(x, y, z)$ is the spectrum of the i 'th Gaussian component in the absorption and one solves for the I polarization leakage $A(y, z)$ and NGAUSS component frequency separations $B_i(y, z)$. Any association of $B_i(y, z)$ with magnetic field values is left to the user.

3.1 Inputs

The inputs for ZAMAN are very similar to those for AGAUS. INNAME *et al.* specify the V polarization cube which must be in transposed form with frequency/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 17. OPTYPE = '2SID' says to solve Equation 3, 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 3. DDOUTPUT 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 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 vanishingly small if ZAMAN is applied to a cube written by ZAMAN with the leakage 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 ZAMAN or from your knowledge of the noise in your data cube.

ZAMAN has one additional adverb, DOCONT, for which the special value -2 causes changed behavior in the fitting. (A special value was chosen to avoid accidental use of this option since DOCONT is used by a number of other tasks.) If the Vpol image cube was produced from continuum-subtracted visibilities (*i.e.*, the output of UVLSF or UVLIN), then the actual Ipol leakage will be proportional to the continuum subtracted Ipol, $I(x, y, z) - C(y, z)$ rather than $I(x, y, z)$. ZAMAN does the fitting differently when this option is selected




Figure 7: First OH spectrum to fit, initial “guess” has Gaussians but the three $B_i(y, z)$ are zero.

and records the fact in the ZE table by setting the keyword ABSORPTN to 2, rather than the usual 1 for normal absorption modeling. The post-fitting plot task XG2PL uses this keyword value to honor the changed meaning of the leakage gain.

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 (DOTV = 1) or, as shown in the present figures, on the TV (DOTV = 2). The first spectrum to be fit (from a model of the absorption shown for AGAUS) is illustrated in Figure 7 showing that the first guess has zero for the three $B_i(y, z)$, 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

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




Figure 8: First OH spectrum to fit, plot after fitting.




Figure 9: Second OH spectrum to fit, initial "guess" has Gaussians and values for $B_i(y, z)$.

menu options in the same way as AGAUS, 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 8. The second pixel starts with the $B_i(y, z)$ of the last fit (Figure 9) and so is 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 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 10. 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 ZAMAN 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, ZAMAN 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 ZAMAN 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; AGAUS establishes which component is which. After the flagging or revisiting, the list is cleared. The first column of the menu includes:

EXIT	Exit AGAUS, writing output images if D00OUTPUT 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 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




Figure 10: ZAMAN fit and image editing screen.

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
SHOW IMAGE F3	Enter image interaction with field of component 3
SHOW IMAGE EF3	Enter image interaction with uncertainty in field of component 3

There is 1 or, if OPTYPE='GAUSS' 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 11, these options are

Figure 11: ZAMAN image of $B_1(y, z)$ 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 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.

When you have finished getting the images just the way you want them, you may write them out as *AIPS*




Figure 12: Example spectrum from a single pixel for fits done by AGAUS and ZAMAN. Data are in black, models in green, and model components in cyan. The model line at left had a width of 8 channels with a splitting of -1 channel, the center line had a width of 15 channels with no splitting, and the line at right had a width of 10 channels and a splitting of $+1.5$ channels. A leakage gain of 0.01 was applied to the full Ipol signal.

image files. Select the SET DOOUTPUT option until its value, shown at the top of the screen, is what you want. In ZAMAN, 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 Post-fit plotting

The images produced by XGAUS, AGAUS, ZEMAN and ZAMAN 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 a new task.

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 AGAUS 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 or ZAMAN 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 12 for the absorption model.

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. For absorption models, CPARM also offers the option of plotting the I spectra as optical depth rather than as observed. 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).

5 Image model creation

AIPS Memo 118 describes a wide variety of modeling tasks which will not be discussed here. A new, simplified task was written to provide direct studies of the fitting in XGAUS, AGAUS, ZEMAN, and ZAMAN.

The new task is called MODAB and has a small number of adverbs. COORDINA, CELLSIZE, APARM(1), and APARM(2) set image header parameters with values that are not otherwise used. IMSIZE and CPARM(3) set the size of the image cube in pixels. APARM(4) through APARM(9) define a Gaussian which is the continuum “source” when imaging in absorption or is used to scale the emission as a function of position when making an emission cube. APARM(10) defines the fractional leakage of the I polarization into V and FLUX sets the noise level in Jy/beam. OPTYPE = 'EMIT' requests an emission cube, otherwise an absorption cube is produced.

A text file, pointed at with INLIST contains up to 99 spectral features to be spread over the cube. Each line must contain

1.	Central optical depth (unitless) or peak brightness (Jy/beam)
2.	Central spectral channel (pixels)
3.	Full width of line at half-maximum (pixels)
4.	Separation of R and L polarization in channels (pixels)




Figure 13: Measured relative widths (s/σ) in green and splittings (d/σ) in red plotted against the model splitting δ/σ .

6 Systematic errors

The image model task described above was used in both emission and absorption with high signal-to-noise to check for “systematic errors” in the answers returned by XGAUS, AGAUS, ZEMAN and ZAMAN. The I polarization Gaussian width, not surprisingly, is measured to be wider than σ_i since the emission profile is widened by the splitting due to the magnetic field. The quality of the fit to a Gaussian also deteriorates as the emission line becomes less Gaussian in shape when the splitting becomes significant with respect to the unsplit line width. The measured splitting in V polarization is computed via Equation 17 and so would appear to suffer from the bias in the measured σ_i squared. Experimentation with the modeling tasks shows, however, that the bias in the splitting is almost the same (fractionally) as in the width. A variety of model separations (δ) were tried at several fixed component widths (σ) with very good signal-to-noise. The apparent Gaussian full-widths at half maximum (s) were measured with XGAUS and the apparent Zeeman splittings (d) were measured with ZEMAN. When plotted as relative widths (s/σ) and relative separations (d/σ) against δ/σ , as in Figure 13, the curves for the tested values of σ were approximately the same. The plotted fits to the curves in Figure 13 are given by

$$s/\sigma = 1.0 - 0.010(\delta/\sigma) + 0.7485(\delta/\sigma)^2 \quad (19)$$

$$d/\sigma = 1.0 - 0.002(\delta/\sigma) + 0.6827(\delta/\sigma)^2 \quad (20)$$

with an rms of the fit of about 0.0002.

Note a certain difficulty with these formulæ. Users of XGAUS, AGAUS, ZEMAN and ZAMAN will obtain values of s and d and will have to find δ and σ by some inverse method. Note too that, while doing these tests,

some rather abnormal values were found, especially for d , despite the rather good signal-to-noise. It may be advisable to use the MODAB modeling task to simulate your data in order to determine the real uncertainties and systematic errors which apply.