

A procedure to make radio/optical composite images

Juan M. Uson (NRAO)

Images of radio sources can be enhanced by combining them with data from other wavelengths, especially so with optical data. This uses the sharpness of the optical image to make the usually diffuse radio emission stand out and not look out-of-focus. It also draws on the familiarity of the viewer with the night sky to provide useful context.

In order to do so, the data should be combined prior to making the three-color image. This step could be avoided by using modern graphics tools that allow layering; but this would require some experiments to find the best software tools and learn how to use them. I describe here the method that I have used within the AIPS package to make images like those included in the figures.

I have used the DSS (POSS-II and Palomar-UKIRT) to obtain near-IR, R and B images. On occasion, the radio source of interest lies in the zone that has been released by the Sloan survey (SDSS) which then provides sharper optical images that are nearly free of defects such as cosmic rays, satellite trails, scratches and saturation of the brightest stars.

One must first choose the area of the desired image in order to obtain the proper size optical images (using the STScI or SDSS servers). After loading the optical images (fits format) the AIPS task 'SKYVE' will place the proper coordinates in the header. The user can choose the epoch of the coordinates as well as definite values for these. In order to obtain the coordinates that were provided within the FITS file, one must set CPARM(1) = 25 and CPARM(5) = 100 (any disallowed value should work).

It is useful at this point to inspect the images and correct defects (for example saturated stars can be replaced by other unsaturated stars appropriately scaled). This can be done using the POPS command interface as long as it does not involve too

large a number of pixels as it is fairly time-consuming and labor-intensive. It can also be done using other analysis packages such as IRAF. Often, the seeing of the near-IR image is significantly smaller than that of the B and R images, which means that the stars will acquire colored halos. In addition, the information contained in the near-IR images is not always distinct. In such cases, I have generated an average 'green' (G) image from the R and B images (suitably scaled to the same units, if necessary).

The radio images should also be clean although this is best done previously at the imaging stage. This can be done with the AIPS routine 'BLANK' (opco 'TVCU' DOINVERS > 0 to blank the interactively set blotch, DOINVERS < 0 otherwise). I have used continuum images and 'total HI images' to make the images described here.

Once the images are cleaned, it is important to ensure that the graphic resolution is sufficient. The images should not reveal the underlying pixels when printed to their intended size. For posters, this should be 18" by 24" (in general). For handouts, covers, the Image Gallery, this should be 8.5" by 11". The prints should have a density of at least 300 dpi. This translates into image sizes of at least 5400 × 7200 pixels for the posters and 2550 × 3300 pixels for the handouts. Most of the starting images are unlikely to have as high a pixel density. Of course, it is possible to make images with small pixels for publication. One only has to run the "final" IMAGR once more with small enough pixels. However, this step is not strictly necessary as the images can be suitably interpolated. It is useful to choose a "reference" image which can be one of the radio images or one of the optical images. This image should be interpolated (for example using OGEOM in AIPS) to the necessary number of pixels (in OGEOM, set APARM(4) to the wanted scaling factor, in the example of Fornax-A shown below I used APARM(4) = 20 to get from the original 512 × 512 image with pixels of 7" to a 10000 × 10000 image with pixels of 0.35" (IMSIZE = 10000 selected the central pixels of the scaled image). Next, use HGEOM to scale the remaining images and re-grid them as well on the same coordinates as those of the "master" image. It is necessary to set the coordinates of the central (reference) pixel in the output image (the same as that of the OGEOM image above), the image size, and the order of the polynomial (I used 7).

Since we have three planes (RGB) available, the images have to be combined. It seems best to add the radio-continuum to the "red" image and the HI to the "blue" image. I have attempted to use the "green" plane as well but the output often seems too surreal. It might be OK to use this plane for molecular emission but I have

not tried it. If the radio emission is only continuum (as in the case of Fornax-A, see below) it seems best to display the radio emission using all three planes to give it a "flaming" look that can be quite spectacular (similar to what can be generated with the AIPS routine TVPHLAME with the red-based hues).

The combination of optical and radio images works because the spatial structure of these images is generally different enough to prevent mutual blurring. However, it is important that the images have "true" zero levels before combination as any "negative" value in one of them will appear as a "ghost" feature in the combined image. I describe here a procedure that I have found to work reasonably well:

- 1) Use IMEAN to obtain a robust value for the sky background level of the optical image from the fit to the histogram of pixel values. It helps to set the average (PIXAVG) to the approximate level on input and an rms (PIXSTD) which is somewhat smaller than the true value (80-90% works well). Iterate until the procedure converges. In order not to display too much noise, set the "zero" level to $n(\text{rms}) \times \text{the average}$; n can be as low as 1.5, but 2 is in general safer. In the case of Fornax-A I used the higher thresholds due to strong gradients in the sky levels of the optical images. For example, the R band image gave a robust average of 3463 ± 388 counts per pixel and I set the zero level to 4250 because of the sky gradient.
- 2) Subtract the zero level (for example with the AIPS routine MATHS, opcode 'POLY', CPARM(1) = "-zero-level", CPARM(2) = 1, CPARM(3) = 0) and set all negative values to zero (for example with AIPS routine BLANK, opcode 'SELC', DPARM(3) > highest value, DPARM(4) = 0, DPARM(5) = 1, DPARM(6) = 0). Blanking with zeroes is necessary to combine the images as described below.
- 3) Decide on the lowest level for the display of the radio data (perhaps $2 \times (\text{rms})?$) and subtract that value from the image (again using MATHS and BLANK to set all pixels below such value to zero). Alternatively, it is useful sometimes to BLANK values below the threshold without subtracting it with MATHS. This can be done simply by setting DPARM(4) to the value of the threshold with the remaining parameters as described in (2). I used DPARM(4) = 0.1 mJy/beam in the example shown below. Thus, the radio data have a "floor" that is above zero which makes even this lowest level noticeable in the composite image.

- 4) Decide on a relative intensity scale between the radio and the optical image. This is best done after displaying them separately using the full three-color scheme.
- 5) For example, run AIPS program TVRGB using the three optical images. Set PIXRA to the limits appropriate to the red image, DPARM(1,2) to those of the green image and DPARM(3,4) to those of the blue image. The initial range can be modified somewhat (decreased) interactively. The program allows displaying a sub-region or skip integer numbers of pixels to make a portion of the image fit on the TV display. A different set of parameters sets the pixels used to make the (PostScript) output file. It is useful to experiment with the "original" images as their smaller size allows for more rapid results for these tests. The test can also be done on the images prior to blanking which might be useful to decide on the offset levels. I enclose below an example from the Fornax-A optical data obtained through TVRGB with inputs: PIXRA = 4250,40000; DPARM = 4150,40000,4000,40000; slightly modified using the interactive cursor (Figure 1). This interactive adjustment operates on the values sent to the display and thus can only narrow the input range. The selected values can be printed to the message file by choosing 'SHOW PIXRANGE' in the interactive menu within the display. Based on this image I decided to use ~34000 for the high values (instead of 40000 which had been my input choice). After subtracting the zero level, this determines a scale of 0 — 30000 for all three colors in the final composite image.
- 6) A "flame" effect will usually enhance the radio emission significantly. This requires making three replicas from the radio image, but the levels have to be determined first. For example, using the AIPS program TVPHLAME interactively yielded the image of the radio emission shown in Figure 2 below. I used the cursor (CURVAL in AIPS) to determine the actual levels of features that should be enhanced. These are the levels at which the flame effect sets in, i.e. the level where the second color is added to the basic red used in this example. In the example, I found the transition levels to be about 2 mJy/beam, 15 mJy/beam, 50 mJy/beam and 115 mJy/beam (the maximum in the image is 123 mJy/beam) The highest level sets everything above it to be white.
- 7) Make three copies of the radio continuum image subtracting the three low thresholds (in the example 0 mJy/beam,

2 mJy/beam and 50 mJy/beam). In the combined images, the red color will start at 0 mJy/beam, the green will be added for values above 2 mJy/beam (the first level targeted for enhancement) and the blue will be added at values above 50 mJy/beam (the third level chosen in the previous step). This step is done again with MATHS, followed by BLANK to set all negative values to 0. Notice that this also changes the range of values of each image, which needs to be remembered in the scaling step discussed next.

- 8) Scale each of the three images appropriately as they are combined with the red, green and blue optical images (using 'COMB', opcode 'SUM'). Given that the full scale chosen in step #5 above was 0 — 30000, these scaling factors were: $2 \times 10^6, 1$ for the addition of the first radio image to the red optical image (this will make the red channel display values of radio emission between 0 and 15 mJy/beam and saturate above this level, the second level determined in step #6 above); $6 \times 10^5, 1$ for the addition of the second radio image to the green image (the green range of the radio emission is 2 mJy/beam to 52 mJy/beam, close to the third transition level found in step #6); and $4.5 \times 10^5, 1$ for the addition of the third radio image to the blue image (which will show radio emission from 50 mJy/beam to 116.7 mJy/beam). The highest level corresponds to that determined in step #6, i.e. everything above 116.7 mJy/beam will be white. The final levels are best determined interactively.
- 9) What one sees on the display is often different from what appears at the printer. Although it is possible to modify the "gamma" functions of the three colors in order to make the printout look closer to what is seen on the display, this step is best left to Pat Smiley of the NRAO graphics department. A few iterations with her will result in a beautiful image.




Figure 1. Three-color optical image. Notice that although most of the light comes through as white the halos around the stars acquire coloring due to the different seeing conditions when the respective plates were obtained.





Figure 2. Three-color image of the radio continuum emission. The color changes determine the levels used to make the three copies needed for combination with the optical images.



NGC1316


Figure 3. Three-color image of the combined radio continuum and optical emission.

Some images can benefit from using non-linear transfer functions. While this can be done "by hand," i.e. modifying the transfer functions with an image manipulation program, I believe that it is best to use a well-defined algorithm that can be expressed in a simple way in a figure caption or in a description that is attached to the image appropriately. A number of algorithms have been used in the past. For example, the 2MASS Atlas of large galaxies uses a "square-root of (1+log)" algorithm to achieve compression of the dynamic range in their images (Jarrett et al. 2003, p. 554).



Arp 299

Figure 4. Composite with HI in blue, optical data from SDSS (linear function).



Arp 299

Figure 5. Composite of Arp 299 with square-root transfer function for the three optical bands. The blank areas were not observed but could be "filled in" with the optical data from the SDSS used in Figure 4. This image has not been

optimized and shows saturation of a few stars, notably the bright one near the center of the lower edge.

I have experimented with this algorithm for the Arp 299 composite image shown above (Figure 5) and found the compression to be excessive. After some tries, I found that a simple square-root compression (for the optical data) led to good results. The HI was "raised" from the zero level by a constant 20 Jy/beam m/sec to allow displaying low-level features but a linear transfer function seemed appropriate. Figure 4 was constructed using linear transfer functions (and SDSS data) while I used deeper optical data with square-root transfer function for Figure 5. Black areas reveal the limited coverage of the deeper data (it would not be difficult to mosaic these deeper images onto the SDSS data that cover the wider field). Other images can benefit from using both (raising the levels by a constant followed by some non-linear transformation (square-root, logarithm). The choice must be guided by what one wants to show.

A non-linear transfer function can also be used in the case of Fornax-A discussed above to enhance the interaction between the two galaxies that is revealed by the low-level connection visible in the optical data. Once again, I have used the "square-root" transfer function on the optical data prior to combining them with the radio images, with the results shown in Figure 6.

Because the optical scale (which was 0-30000 for all three images has now been modified by the "square-root," the scaling used to combine the radio and optical images had to be modified as well. Indeed, the "square-root optical images" have been scaled by the factor 150 (a bit smaller than square-root of 30000 in order to allow the central jet to show -once a pixel becomes "white" no amount added to the red or other gun will modify its color) prior to adding them to the radio images as described above.

So, step #8 in the description above would be replaced by:

8) Replace each of the values in the three optical images by their square-root using MATHS (OPCO: 'POWR', CPARM= 0,1,1,0,0.5). Scale each of the three images appropriately as they are combined with the red, green and blue optical images. In my example, these scaling factors were: (2×10^6 , 350) for the addition of the first radio image to the red optical image; (6×10^5 , 350) for the addition of the second one to the green image; and (4.5×10^5 , 350) for the addition of the third radio image to the blue image.



NGC1316

Figure 6. Three-color image of the combined radio continuum and optical emission with square-root transfer function for the optical data which enhances the connection between NGC1316 and its neighbor to the north.

Acknowledgements: I thank John Hibbard for much feedback on the production of these images as well as suggestions on data visualization. Eric Greisen modified a number of AIPS routines to allow displaying and making hard copies of images with an ever increasing number of pixels. Alan Bridle provided valuable suggestions when images looked faded and Pat Smiley contributed her unique sense of what makes an image beautiful.

Reference:

Jarrett, T. H. et al. 2003, AJ. 125, 525.