

Making the Invisible Visible: Converting Multi-band Astronomical Data into RGB Images

Ruining Xin

*The Experimental High School Attached to Beijing Normal University, Beijing, China
18090919870@163.com*

Abstract. Modern astronomy relies on multi-wavelength observations across the electromagnetic spectrum, yet most of these data lie outside the optical regime. Since different physical processes emit radiation at characteristic energies, no single band can capture the full picture of a complex astronomical environment. We introduce a Python pipeline for converting multi-wavelength astronomical observations into scientifically representative pseudo-color composite images. The workflow incorporates FITS ingestion, statistical profiling, chromatic response modeling, and energy-ordered RGB projection. Each step addresses a specific challenge: reading heterogeneous data formats, characterising dynamic range, modelling human colour response, and projecting multiple bands into a perceptually meaningful image. It is found that the arcsinh stretch achieves optimal contrast across high dynamic ranges, compared with four intensity scaling laws: linear, logarithmic, arcsinh, and power-law, on Chandra observations of Cassiopeia A and the Pinwheel Galaxy (M101). This is because arcsinh behaves linearly at low intensities and logarithmically at high intensities, preserving both faint and bright structures within a single display. Additionally, percentile clipping and three-channel color assignment allow spatially overlapping emission structures to be distinguished visually. This lightweight framework provides a transparent tool for generating scientifically meaningful astronomical illustrations.

Keywords: FITS data, astronomical image processing, pseudo-color composite, cone-cell response, multi-band imaging

1. Introduction

Instead of inspecting the universe through a single lens, modern astronomy relies on multi-wavelength observations across the entire electromagnetic spectrum, from long-wavelength radio waves to high-energy gamma rays. Since different physical processes emit radiation at specific energies, we can find information on the structure, temperature and dynamics via observing an astronomical target across multiple spectral bands. Meanwhile, by integrating the observing data from radio, infrared, optical, ultraviolet, X-ray, and gamma-ray, we can obtain clearer images of complex astrophysical environments such as star-forming regions, active galactic nuclei, and supernova remnants.

Human's exploration on the optical band began four centuries ago. In 1610, Galileo pointed a crude refracting telescope at Jupiter and found four moons that had never been recorded before. He not only overturned the old belief that all celestial bodies revolved around the earth, but also opened the era of systematic observation of the universe with instruments.

Decades later, Kepler distilled his three laws of planetary motion from the observational data left by Tycho, Newton explained the reasons behind these laws with his law of universal gravitation and elucidated the laws of light dispersion, and Huygens improved the design of refracting telescopes and discovered Saturn's largest moon, Titan. Newton's reflecting telescope, invented to overcome the chromatic aberration of lenses, established a design tradition that continues to this day.

However, the optical channel, roughly 400 to 700 nanometers, is only one slice of the electromagnetic spectrum; it is just one narrow slice of this continuum. The twentieth century saw astronomy expanding successively into radio, infrared, ultraviolet, X-ray, and gamma-ray windows, each opened by a new generation of detectors and observatories. Modern facilities such as the James Webb Space Telescope (JWST) [1] observe the infrared universe, the Atacama Large Millimeter Array probes cold molecular gas, the Chandra X-ray Observatory maps hot plasma around supernova remnants and active galactic nuclei, and the Fermi Gamma-ray Space Telescope surveys the high-energy sky. The Event Horizon Telescope has even produced images of the shadows of supermassive black holes at the centres of the galaxy M87 and our own Milky Way [2, 3]. Together these observatories now observe the same cosmic objects in dozens of non-visible bands.

This paper describes a six-step pipeline that performs such a conversion. The pipeline was developed upon a compact Python library that wraps `astropy.io.fits` for data reading and `matplotlib` for visualisation. The six steps are: read FITS data, compute pixel statistics, model the human cone-cell response, convert each wavelength to a colour triplet, project multiple bands into a single RGB image, and render the final composite. The procedure is illustrated on observations of two well-known objects: M101 (the Pinwheel Galaxy), observed in four bands, and Cassiopeia A, observed in three X-ray energy bands. The novelty of the present work is not the individual components, which are standard practice in the astronomical community [4], but the explicit documentation of the entire pipeline on real data.

2. Methods

2.1. Step 1: reading FITS data

The Flexible Image Transport System (FITS) is the standard astronomical data format, maintained by NASA and ESA [5]. The Python library `astropy` [6] exposes a FITS file through a `fits.HDUList` object. The function `load_fits(path)` opens the file, extracts the primary data array, casts it to a 64-bit float, and replaces any NaN or Inf values with zero.

2.2. Step 2: pixel statistics

Astronomical images are non-negative arrays whose values are event counts, photon rates, or surface brightnesses. Their distribution is highly skewed: a few bright sources sit on a broad background of barely-detectable photons. The function `image_stats` reports the shape, mean, median, standard deviation, the 1st and 99th percentiles, and the dynamic range. For the X-ray image of M101 used in this study, the dynamic range exceeds 10^4 , which already explains why a linear stretch cannot show faint and bright features simultaneously.

2.3. Step 3: cone-cell response

The human retina contains three classes of cone photoreceptor. The standard approximation uses three Gaussian response curves centred near 650 nm (L-cones, "red"), 550 nm (M-cones, "green"), and 450 nm (S-cones, "blue"), with a width parameter $\sigma \approx 50$ nm that matches psychophysical measurements [7]. The function `cone_response(wavelength, sigma=50)` returns the three response values R, G, B for any wavelength in nanometres. These three Gaussian curves are the perceptual basis for converting any non-visible wavelength into a colour, because every visible hue corresponds to a particular triplet of cone responses.

2.4. Step 4: wavelength to RGB

The function `wavelength_to_rgb` maps a visible wavelength to a discrete RGB triplet in the range [0, 255] using a piecewise approximation of the CIE colour-matching functions: red rises from 380 to 440 nm, vanishes between 440 and 580 nm, and rises again from 580 to 650 nm; green rises between 440 and 490 nm and falls between 580 and 650 nm; blue falls from 440 to 490 nm. A linear intensity falloff at the spectral extremes compensates for the steep decline of the standard luminosity function. Wavelengths outside the visible range return black. Each branch corresponds to a physical intuition.

2.5. Step 5: intensity scaling

A raw X-ray or infrared pixel value is not a luminance. The human perception of brightness is roughly logarithmic, so a linear stretch leaves the faint sky invisible and a log stretch washes out bright details. The pipeline implements four scaling laws, linear, logarithmic, arcsinh, and square-root, through the `scale_data` function, each mapping a pixel value to the [0, 1] interval after percentile clipping with `auto_vmin_vmax`. The arcsinh stretch, $I_{\text{out}} = \text{asinh}(I_{\text{in}}/\beta) / \text{asinh}(v_{\text{max}}/\beta)$ with a softening parameter $\beta = 0.1 v_{\text{max}}$, behaves linearly at low intensity and logarithmically at high intensity. The technique was introduced by Lupton et al. [4] for SDSS colour composites and is widely used because it preserves colour balance across a wide dynamic range.

2.6. Step 6: pseudo-color projection

The function `color_project(data, wavelength)` scales the data and multiplies the resulting [0, 1] array by the cone-response triplet for the given wavelength. The result is an RGB image whose hue intuitively corresponds to the observation band: infrared is rendered in red, ultraviolet in blue, radio in red or purple. For multi-band composites, the function `make_rgb_composite` takes a list of band dicts (each containing the data, a central wavelength, and an explicit R/G/B channel assignment) and adds the projected contribution of each band into the corresponding channel. After summation, the composite is renormalised to [0, 1] and rendered with `matplotlib.pyplot.imshow`.

3. Results

As shown in Figure 1(a), the panel displays electromagnetic spectrum from radio to gamma-ray frequencies with major observatories marked in individual wavelength, in which, visible light occupies only a narrow slice of the entire spectrum.

Figure 1 (b) shows the normalised Gaussian cone-cell response curves centred at 450 nm (S-cones, blue), 550 nm (M-cones, green), and 650 nm (L-cones, red) with $\sigma \approx 50$ nm [7]. The relative activation of these three cone types provides the primary basis for human chromatic encoding.




Figure 1. Electromagnetic spectrum and human cone-cell response

(a) The electromagnetic spectrum from radio to gamma-ray frequencies, with representative observatories annotated for each band

(b) Normalised cone-cell response curves for S-, M-, and L-cones

Figure 2 shows the histogram of the central 1200×1200 pixel crop of the Chandra X-ray image of M101 is dominated by low-count pixels, with a long tail extending to bright point sources. The 1st, 50th, and 99th percentiles are marked, illustrating the highly non-Gaussian nature of the data. The dynamic range between P1 and P99 exceeds 10^4 , which is far beyond what a linear display stretch can render.




Figure 2. Pixel-value histogram of the M101 X-ray observation

Each panel in Figure 3 shows the central quarter of M101 in a different band, projected onto the human RGB space using the cone-response model. The X-ray panel (Chandra) shows unresolved point sources and diffuse hot gas; the UV panel (GALEX) highlights star-forming knots in the spiral

arms; the optical panel (Hubble) reveals dust lanes and the older stellar disk; the infrared panel (Spitzer) emphasises warm dust heated by embedded star formation. Because each panel is arcsinh-scaled with its own percentile limits, structural details spanning different dynamic ranges are visible in every band.




Figure 3. Single-band colour projection of M101

In Figure 4, linear scaling leaves the remnant faint and saturates the brightest knots; logarithmic scaling brightens the faint outer filaments but reduces the contrast of the bright ring; the square-root stretch improves contrast modestly; the arcsinh stretch provides the most balanced result, with faint wisps, the bright main shell, and small-scale filamentary structure all visible simultaneously. The arcsinh stretch is the recommended default for X-ray supernova-remnant images because it preserves structure across a large dynamic range.




Figure 4. Comparison of four intensity scaling methods on Cas A (1.5–3.0 keV X-ray)

In Figure 5, the three leftmost panels show the individual X-ray energy bands (0.5–1.5 keV, 1.5–3.0 keV, 4.0–6.0 keV) in grayscale, all arcsinh-scaled. The rightmost panel shows the pseudo-color composite, in which the soft 0.5–1.5 keV band is encoded in red, the medium 1.5–3.0 keV band in green, and the hard 4.0–6.0 keV band in blue. Soft emission dominates the outer filaments, medium emission fills the main shell, and hard emission traces the fragmented inner ring. The composite conveys the energy structure of the remnant at a glance and explains why multi-band colour rendering is now standard in high-energy astrophysics [8].




Figure 5. Multi-band pseudo-color composite of Cassiopeia A

In Figure 6(a), it displays four scaling functions plotted in normalised coordinates. The linear curve passes through the origin and peak equally; the logarithmic curve suppresses bright values; the arcsinh curve ($\beta = 0.1$) tracks the linear curve at low intensity and the logarithmic curve at high intensity; the square-root curve lies between them. In Figure 6(b), it shows histogram of the Cas A pixels after linear and arcsinh scaling. Linear scaling concentrates most pixels near black, while arcsinh redistributes them across the $[0, 1]$ display range, explaining the improved visibility of faint structure in the corresponding image.




Figure 6. Mapping curves of four intensity scaling laws and the resulting pixel distributions

4. Discussion

The pipeline presented here is of an instructional and communication tool that makes the choices of the rendering process explicit. The most subtle of these choices is the intensity scaling function. The four scaling laws implemented in `scale_data` were compared side-by-side on the same X-ray data (Figure 4). The arcsinh stretch performed best on the Cas A data, because the supernova remnant contains both bright hotspots comparable to the saturation of the Chandra detector and faint filaments only a few counts above the background. The linear stretch buried the faint filaments; the logarithmic stretch washed out the hotspots; only the arcsinh stretch preserved both.

The percentile-based default limits (`auto_vmin_vmax` with the 1st and 99th percentiles) are a robust starting point but are not always optimal. In the M101 X-ray image, the 99th percentile is

dominated by a small number of bright X-ray binaries, which forces the diffuse hot gas to appear faint. Adjusting the upper limit to a higher percentile recovers the diffuse component. This observation aligns with the broader recommendation of Rector et al. [9] that intensity limits be chosen deliberately for each scientific case, not blindly.

The cone-response model is also a simplification. The human visual system does not compute a hue from three Gaussians; it performs a complex non-linear transform that includes opponent channels, spatial masking, and adaptation [7]. Despite this, the Gaussian approximation captures the essential property used here: any visible wavelength can be associated with a reproducible RGB triplet, and any non-visible wavelength can be associated with a designated primary colour. This is the basis on which all modern multi-band astronomical images are produced, whether they are JWST press releases [10], Chandra supernova-remnant montages [8], or Event Horizon Telescope reconstructions of black-hole shadows [2, 3].

A practical caveat is that the pipeline operates on a single two-dimensional FITS image per band. Instruments such as JWST or the upcoming Nancy Grace Roman Space Telescope produce many exposures per band, requiring stacking, drizzling, and mosaic assembly before the present pipeline can be applied. The piecewise wavelength-to-RGB function is also a rough approximation; for visualisation purposes more accurate conversions using CIE 1931 XYZ functions are available, but the educational value of the present approximation, in which each branch corresponds to a physical intuition, is preserved.

5. Conclusion

A six-step pipeline has been implemented, tested, and documented for converting wide-band astronomical FITS data into visually meaningful images. The pipeline reads FITS data, computes pixel statistics, models the human cone-cell response, converts wavelengths to RGB, scales intensities with a choice of four laws, and renders single-band or multi-band pseudo-color composites. The pipeline has been demonstrated on real observations of the Pinwheel Galaxy (M101) and Cassiopeia A (Cas A), provided by the Chandra X-ray Center openFITS archive. The results show that the arcsinh intensity scaling function, combined with percentile-based display limits, produces the most informative rendering of X-ray data, and that the band-to-channel mapping, R = soft, G = medium, B = hard, is a defensible default for visualising the energy structure of supernova remnants. The complete code is short, uses only numpy, astropy, matplotlib, and scipy.

References

- [1] Gardner, J. P., Mather, J. C., Clampin, M., et al. (2023). The James Webb Space Telescope mission. *Publications of the Astronomical Society of the Pacific*, 135(1048), 068001.
- [2] Event Horizon Telescope Collaboration (Akiyama, K., et al.). (2019). First M87 Event Horizon Telescope results. I. The shadow of the supermassive black hole. *The Astrophysical Journal Letters*, 875(1), L1.
- [3] Event Horizon Telescope Collaboration (Akiyama, K., et al.). (2022). First Sagittarius A* Event Horizon Telescope results. I. The shadow of the supermassive black hole in the center of the Milky Way. *The Astrophysical Journal Letters*, 930(2), L12.
- [4] Lupton, R., Blanton, M. R., Fekete, G., et al. (2004). Preparing red-green-blue images from CCD data. *Publications of the Astronomical Society of the Pacific*, 116(816), 133–137.
- [5] Pence, W. D., Chiappetti, L., Page, C. G., et al. (2010). Definition of the Flexible Image Transport System (FITS). *Astronomy & Astrophysics*, 524, A42.
- [6] Astropy Collaboration (Price-Whelan, A. M., et al.). (2022). The Astropy Project: Sustaining and growing a community-oriented open-source project and the latest major release (v5.0) of the core package. *The Astrophysical Journal*, 935(2), 167.

- [7] Stockman, A., & Sharpe, L. T. (2000). The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research*, 40(13), 1711–1737.
- [8] Dubner, G., Castelletti, G., Kargaltsev, O., et al. (2017). Morphological properties of the Crab Nebula: A detailed multiwavelength study based on new VLA, HST, Chandra, and XMM-Newton images. *The Astrophysical Journal*, 840(2), 82.
- [9] Rector, T. A., Levay, Z. G., Frattare, L. M., Arcand, K. K., & Watzke, M. (2017). The aesthetics of astrophysics: How to make appealing color-composite images that convey the science. *Publications of the Astronomical Society of the Pacific*, 129(975), 058007.
- [10] Cutler, S. E., Whitaker, K. E., Weaver, J. R., et al. (2024). Two distinct classes of quiescent galaxies at cosmic noon revealed by JWST PRIMER and UNCOVER. *The Astrophysical Journal Letters*, 967(2), L23.