

# Pretty astronomical pictures

- ❑ Three-color image composition
- ❑ The most important things to know from Photoshop

Good references:

Adam Block 2013, *Deep sky imaging: workflow 1*

Johannes Schedler 2013, *Deep sky imaging: workflow 2*

Both in *Lessons from the masters*, ed. Robert Gendler (New York: Springer)



M 42 from Mees, L-[S II]-H $\alpha$ -[O III]; L from 22 February, spectral lines from 27 March. Semi fashionable handbag from sow's ear.

# New course logistics

- ❑ Remember that if you would like to receive a letter grade for this course, you need to ask by 29 April: [click here to do so](#).
- ❑ There are now **drylab** options for each project, in view of the poor long-range weather forecast and the impending end of the semester.
- ❑ Your completed **term project** is due at the end of the last day of finals week: 10 May 2020.
- ❑ **First drafts**, for those who are taking AST 244W, are due at the end of the day on 3 May 2020.



# LRGB

If one is after pretty pictures, one has already invested a lot more time in L than in RGB, according to the prescription promoted in these lessons.

- ❑ Even in the same time, L images have much higher S/N than R, G, or B, and so will look much more impressive.

$$\left(\frac{S}{N}\right) \propto \sqrt{P_s \Delta t} \text{ if background limited}$$

- ❑ By the same token one would have to commit prohibitively-large (factors of 3-9 larger) exposure times for R, G or B to catch up in S/N with L.
- ❑ Thus ground-based, broadband pretty pictures are almost always LRGB composites.



Today's example: M 101 from Mees, L, R, G, and B, 14 June 2016. This is L. Notice the many artifacts – mostly flat field errors – that were not corrected in calibration.

---

## LRGB (continued)

So one usually uses the multiple-image-plane and Luminosity features of Photoshop to combine L and RGB images, using the RGB as the background, shining through the L image.

- ❑ L acts as a transmission mask, thus suppressing noise in the blank sky in the RGB image, and imposing sharp detail on the bright parts.
- ❑ Enables use of even more subterfuge on RGB, such as coarser binning (which is already part of our recipe) and further blurring within Photoshop to reduce color noise, without changing the resolution of the final product.
- ❑ Penalty: total loss of color and surface-brightness accuracy.
  - Colors are smeared over more L PSFs than they deserve to be.
  - If the L image has been deconvolved too, well ... see last lesson.
  - Now it's more a work of art. Which is of course not all bad.

---

# Color composition

CCDStack is a good tool for combining multicolor Mees images into an RGB color composite. Here are the steps.

- ❑ If your L images are binned 1x1 and your colors 2x2, according to our recommendations in Lesson 2, then load L, and then the colors with L as the reference image, and register (align) the others to it. This will result in an interpolation of the colors to 1x1 binning, as well as lining up all the stars. Then close the L image.
- ❑ Go to [Color Combination](#), specify the images to combine if necessary, and enter the scaling factors.
  - Choose which images you want represented as R, G, and B.
  - Choose [scale factors](#) by which to multiply the images, according to how the colors bring out the structures you want to emphasize.
  - “True color”, to make A0V look white in images with the same exposure time: R:G:B = 0.67:0.80:1.0. Or to make G2V stars look white: R:G:B = 0.50:0.67:1.0.

---

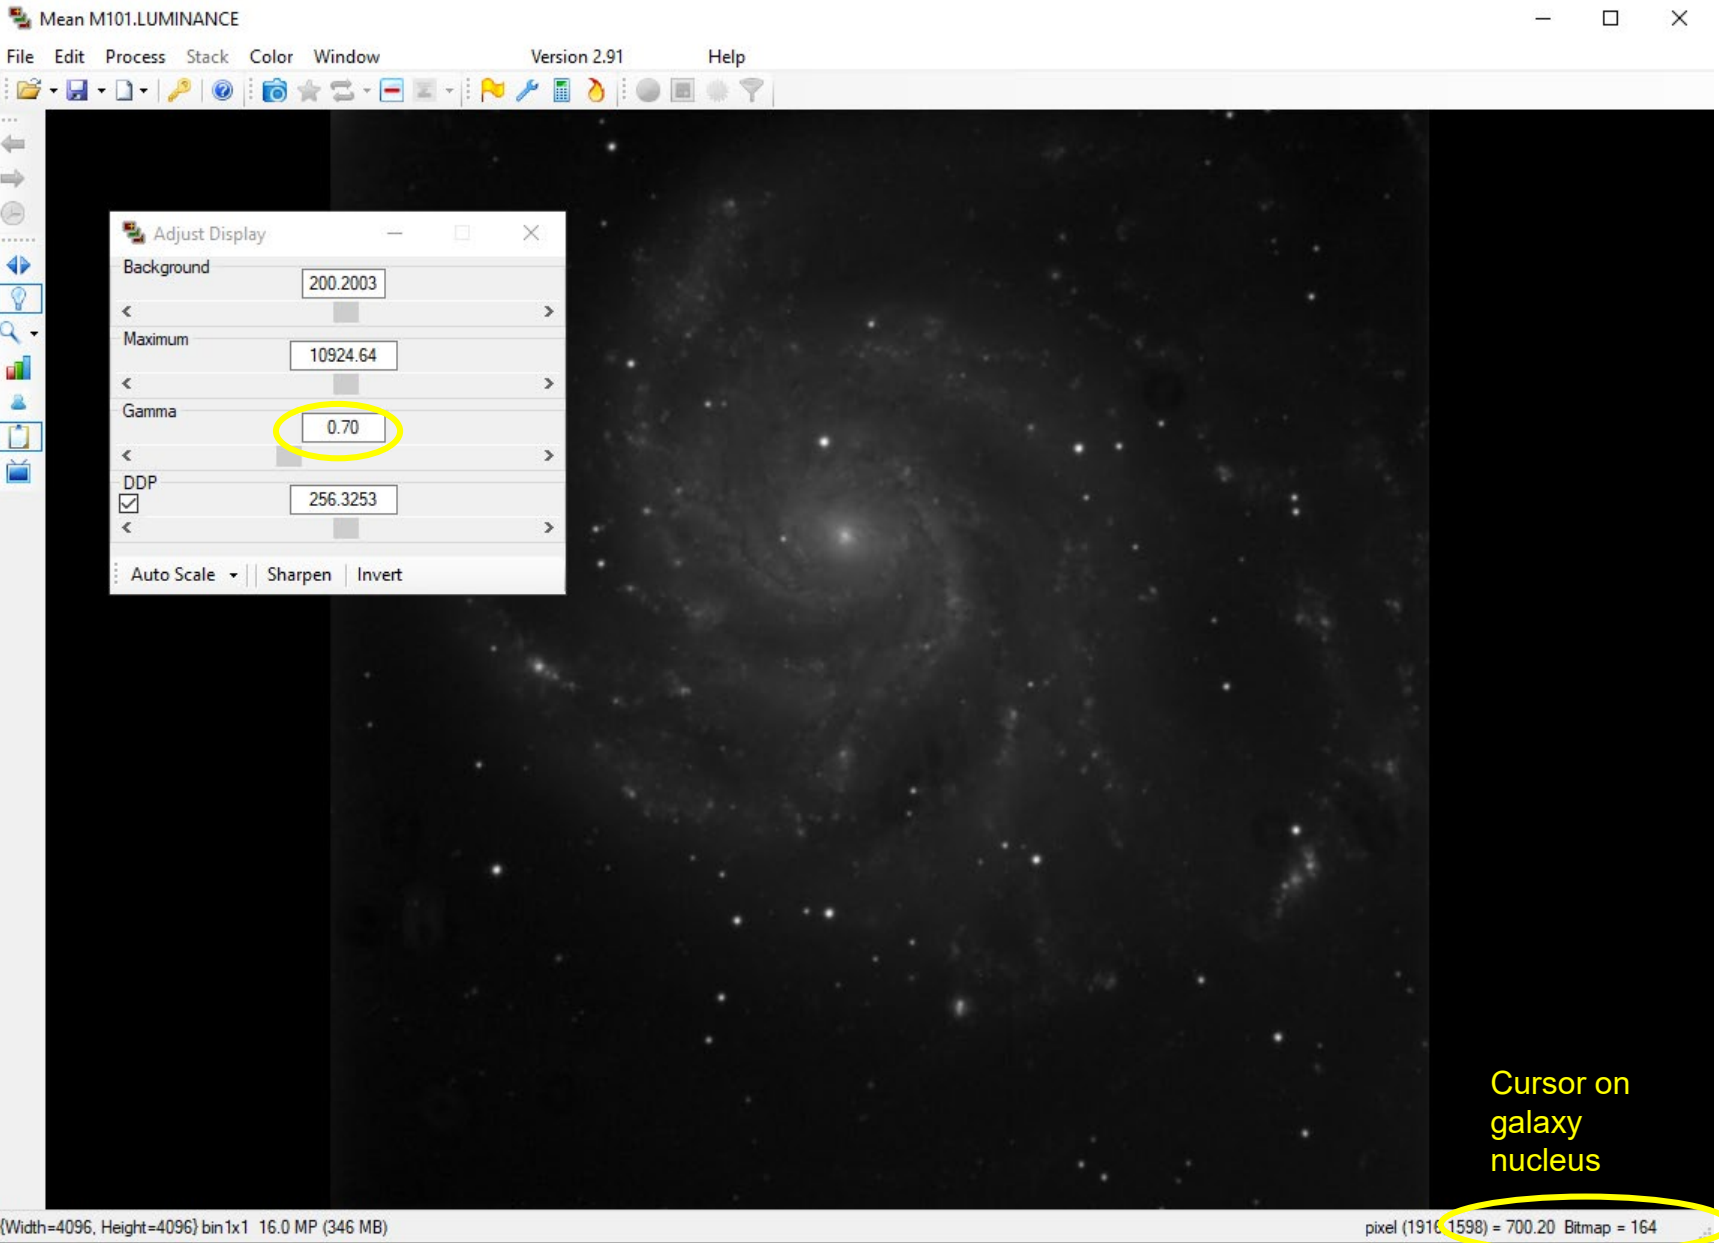
## Color composition (continued)

- ❑ Then click [Apply](#) and you will get an RGB composite.
- ❑ Reload your L images.
- ❑ Now, counterintuitively, you'll need to make them both a bit duller, so that Photoshop can change the component brightnesses either up or down:
  - **Stretch** the display to a lower gamma and maximum, so that nothing important is close to 256 in Bitmap units.
  - [Save scaled data](#) to produce TIFF images.



M101, RGB, full frame,  
DDP stretch





---

## Aside: stretching images

CCDs have DN linearly proportional to power collected; each pixel's signal is sent out as a 16-bit floating point number, and stored in the computer as a 32-bit number. None of this matches displays very well.

- ❑ Computer monitors only display eight bits of brightness (0-256), for the very good reason that eyes can only resolve that many shades of gray. Printing on paper is similar.
- ❑ Thus one must **stretch** (or compress, really) the huge range of astronomical brightness within an eight-bit range for display.
- ❑ Photographic emulsion has a useful compressive feature built in to it, called **reciprocity failure**:
  - Hypersensitized emulsion produces image density linear in power at low light levels, but the response becomes more like logarithmic with brighter lights.



---

## Stretching images (continued)

Each of the image analysis programs we use has a variety of ways to stretch images, notably these:

- ❑ **Gamma.** This maps power exponentially into signal:

$$DN_{\text{display}} = A(DN_{\text{image}})^{\gamma} + B$$

$\gamma = 1$  is of course the same as the original image, but  $\gamma < 1$  compresses the display brightness into a smaller range.  $A$  (“brightness”) and  $B$  (“background”) can be set separately.

- ❑ **Logarithmic stretch:**

$$DN_{\text{display}} = A \log(DN_{\text{image}}) + B.$$

An attempt to mimic what the eye itself does, but it has obvious problems if it's possible for the pixels to have DN values of zero.

---

## Stretching images (continued)

- ❑ **Digital development process (DDP).** This is an algorithm, developed by a professional physicist who happened to be an amateur astronomer, which is meant to mimic the useful features of reciprocity failure. The algorithm is iterative and replicates the stages one would see in an exposed photographic plate immersed in developer.
  
- ❑ **Arcsinh stretch.** At large values,  $\operatorname{arcsinh}(x)$  converges to  $\ln(x)$ ; at small values it's linear (like  $\ln(x)$ ) but  $\operatorname{arcsinh}(0) = 0$  (unlike  $\ln(x)$ ). So it satisfies all the constraints and embodies several useful properties.
  - It also behaves very much like reciprocity failure in hypersensitized emulsion. The DDP stretch is very similar to the arcsinh stretch, as was first noticed long before the invention of DDP by fitting functions to the response of “analog” developed emulsion: Bunsen & Roscoe 1862, Schwarzschild 1899, Kron 1913.

---

## Stretching images (continued)



Linear



DDP

---

## Stretching images (continued)



arcsinh



DDP

---

## LRGB (continued)

Now, on to Photoshop.

- ❑ Ask Photoshop to open ([File > Open](#)) the scaled L and RGB TIFF images you want to use.
- ❑ Choose the tab with the L image, select the whole image (Ctrl+A) and copy (Ctrl+C)
- ❑ Choose the tab with RGB, and paste (Ctrl+V). This creates a layer containing L, above the RGB background.



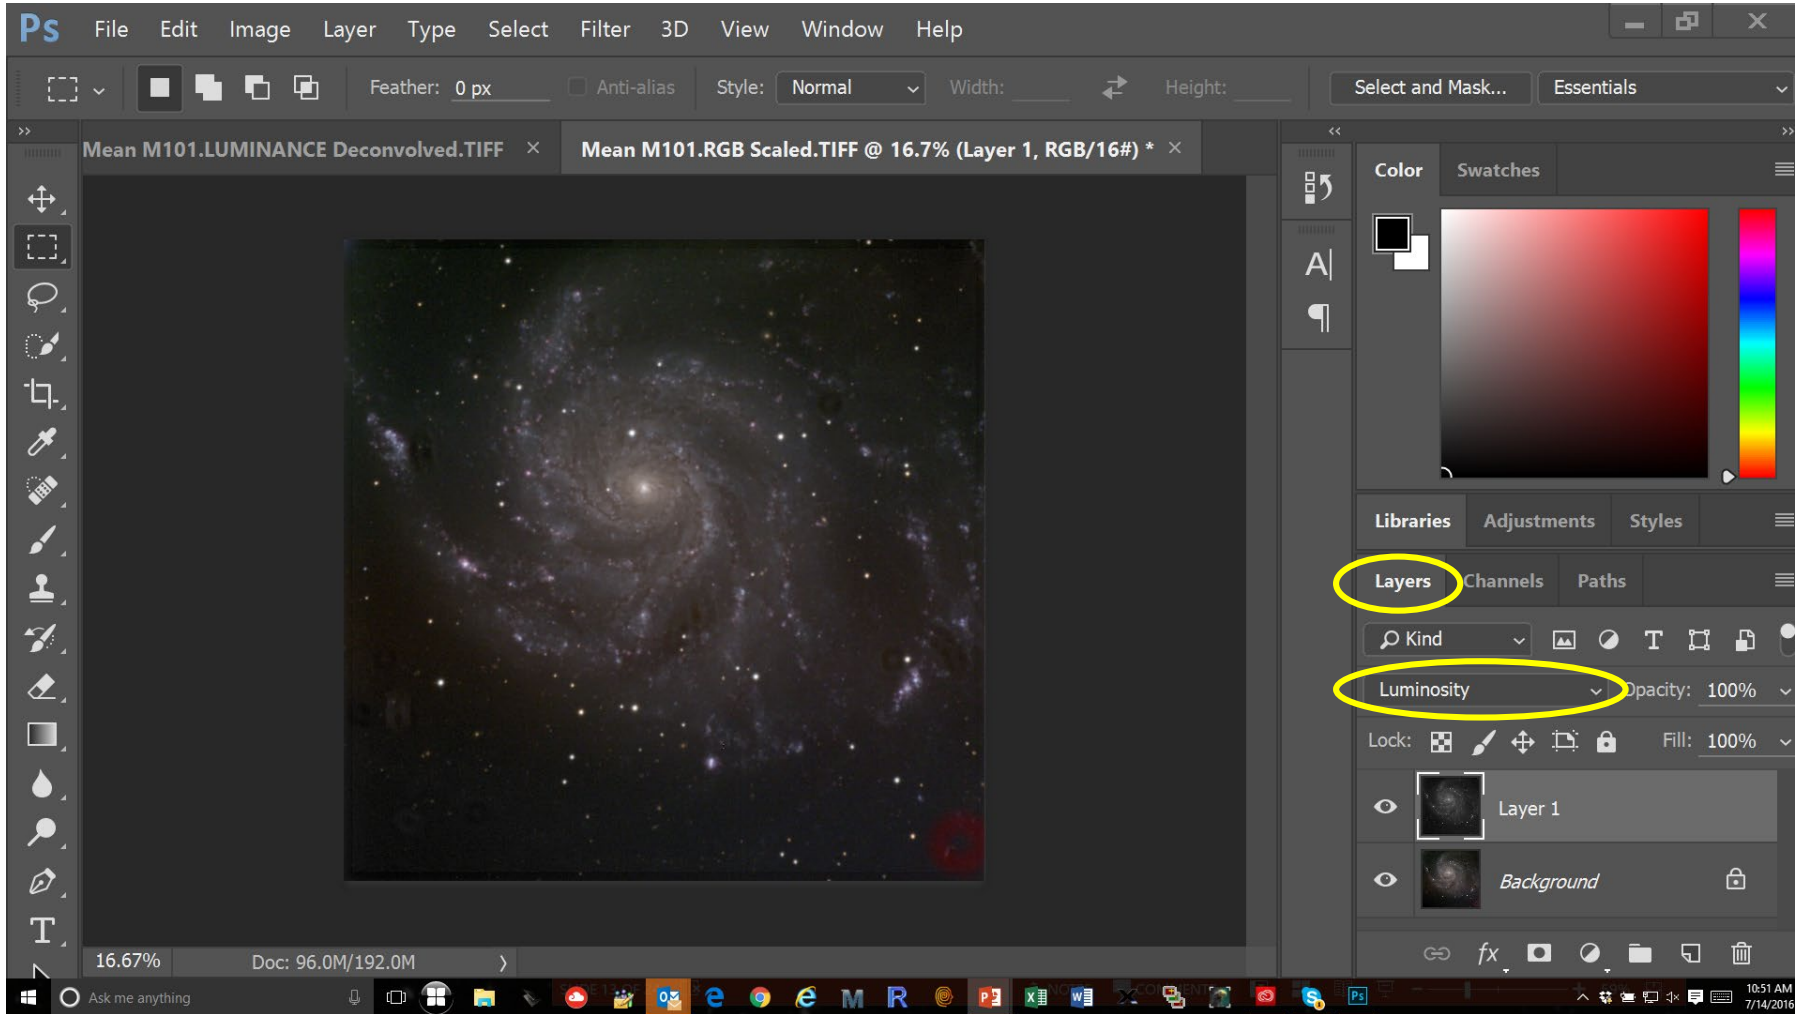
---

## LRGB (continued)

- ❑ You will see this superposition indicated on the [Layers](#) tab at lower right of the Photoshop window.
  - Rename the layers, if you want, by double-clicking the name and typing a new one.
- ❑ Select the upper layer (L, Layer 1 if you didn't rename it), click the dropdown box immediately above it, which starts off saying [Normal](#), and choose [Luminosity](#) (last choice in the list).
- ❑ Now the colors are shining through the L image, and you're ready to start tuning the image to look nice.

Tip: take note of the menu items [Edit > Undo](#), [Edit > Step Backward](#) and [Edit > Step Forward](#), by which you may recover from almost any misstep in Photoshop, or toggle between two steps in the process.

# LRGB (continued)





---

# Tuning up the LRGB image

As long as L and RGB are two separate layers, one may – and should – operate on them separately, but one should always be looking through L onto RGB with the Luminosity overlay, as on the last page.

Here are the tools to use, and *roughly* the order in which to use them

- ❑ [View > Zoom In](#), [View > Zoom Out](#), [View > Fit on Screen](#), invoked from the keyboard with Ctrl+ +, Ctrl+ -, and Ctrl+ 0. You'll use these a lot.
- ❑ Select RGB (background) layer and do [Filter > Blur > Gaussian Blur](#). This essentially convolves the RGB layer with a Gaussian, which smooths out the noise in this layer; it also blurs that layer, but not the final image, as the L layer determines the detail.
  - Do this to suppress “color noise” in the Highlight parts of the image: those parts where L is bright.

---

# Tuning up the LRGB image (continued)

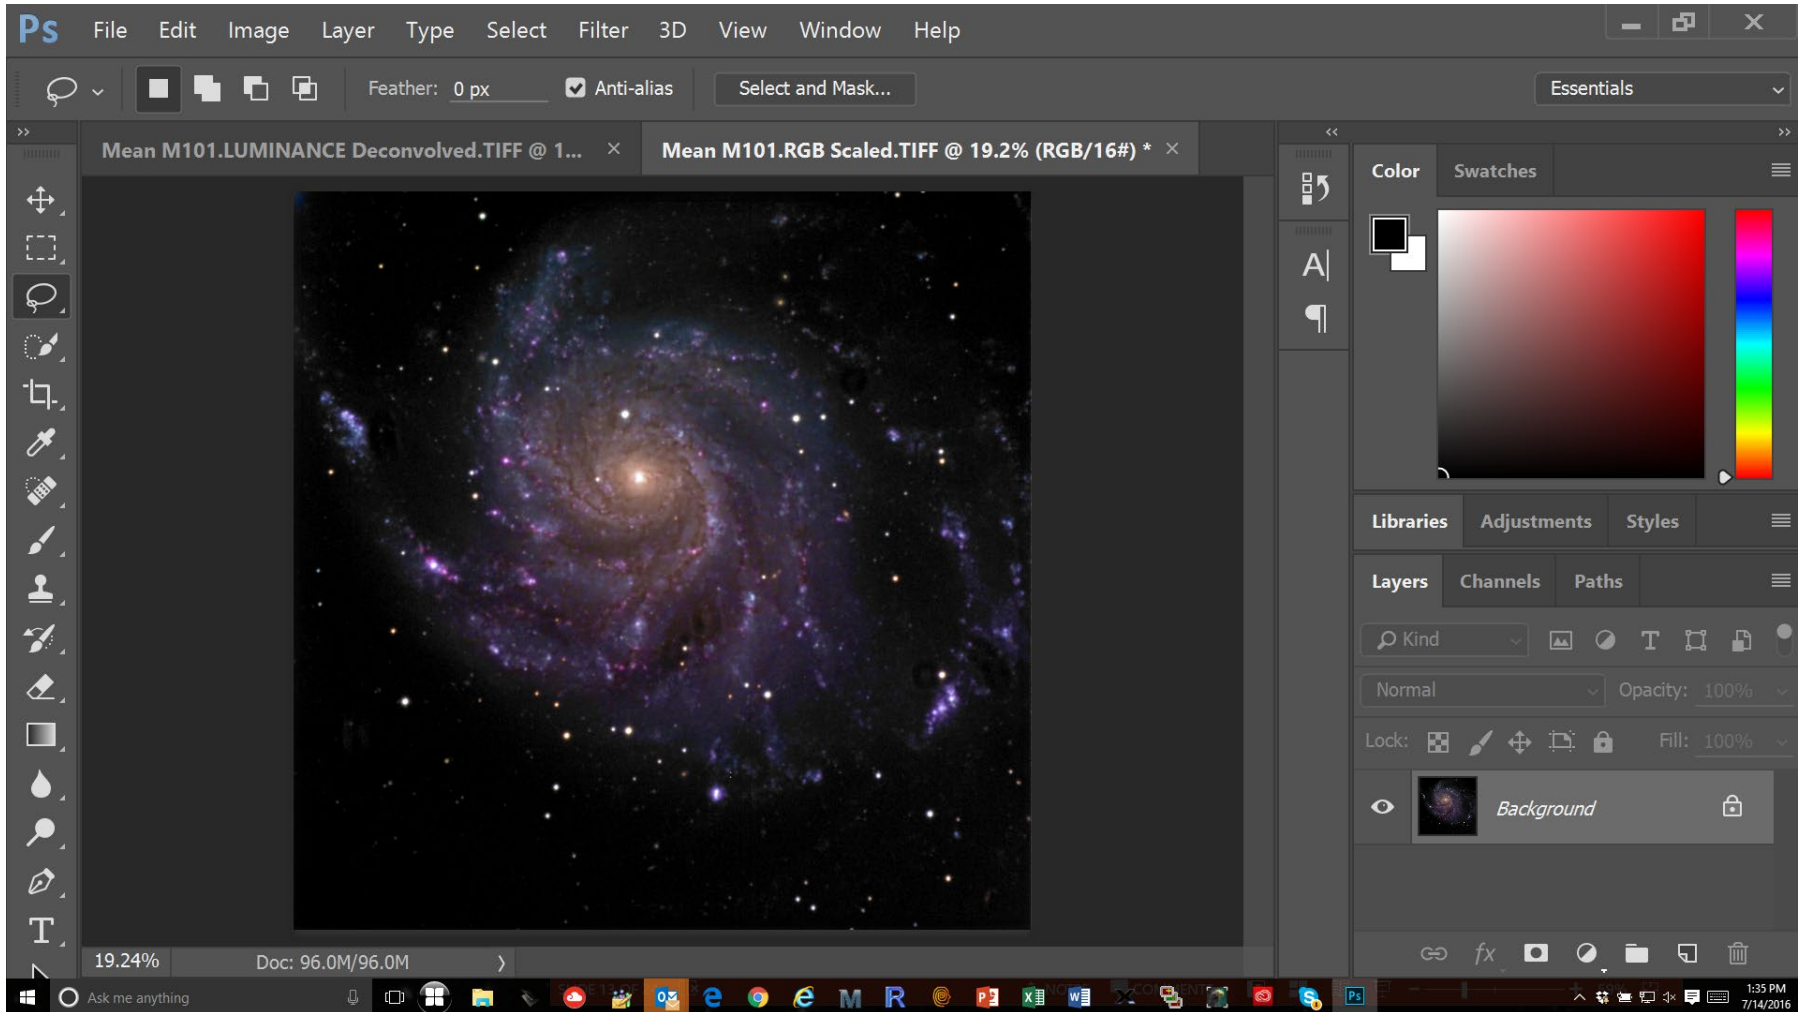
- ❑ Several tools under [Image > Adjustments](#), still just on the RGB layer:
  - [Shadows/Highlights](#). Favorite of Adam Block's; I can't make it work as well as he does, but it is a versatile tool for adjusting signal and color of background and highlights with respect to each other. For example: taking color out of the background without affecting the color of the highlights very much.
  - [Levels](#). Can adjust the black level (left slider), white level (right) or midtone (middle) of RGB all at once or one color at a time.
    - Note especially that raising the black level of RGB – basically cutting off all the signals below the level indicated by the left slider – increases the color saturation of the highlights dramatically.
  - [Curves](#). Enhance or suppress some ranges of brightness with respect to others, in a smooth and nonlinear fashion

---

## Tuning up the LRGB image (continued)



- **Vibrance.** Brightness of color, which is somewhat different from color saturation. Often turned up way too far in APOD images; careful of what it does around stars and compact parts of the highlights. Also has a Saturation slider.
  - **Hue/saturation.** Allows tweaking of colors in mixed ways, like CYM instead of RGB. Adjusting M (magenta) is often useful in making H II regions in galaxies look a bit more like the real thing.
  - **Color balance.** Some of the same features as **Hue**, but allows adjustment separately for shadows, midtones and highlights, and does so in energy-conserving fashion.
  - **Brightness/contrast:** like **Levels** but cruder. Might not help much if you did all the others first.
- 
- After RGB, select the L layer and revisit **Levels** and **Curves**.
  - Important: **raise the L black level last.**

# Tuned-up LRGB image

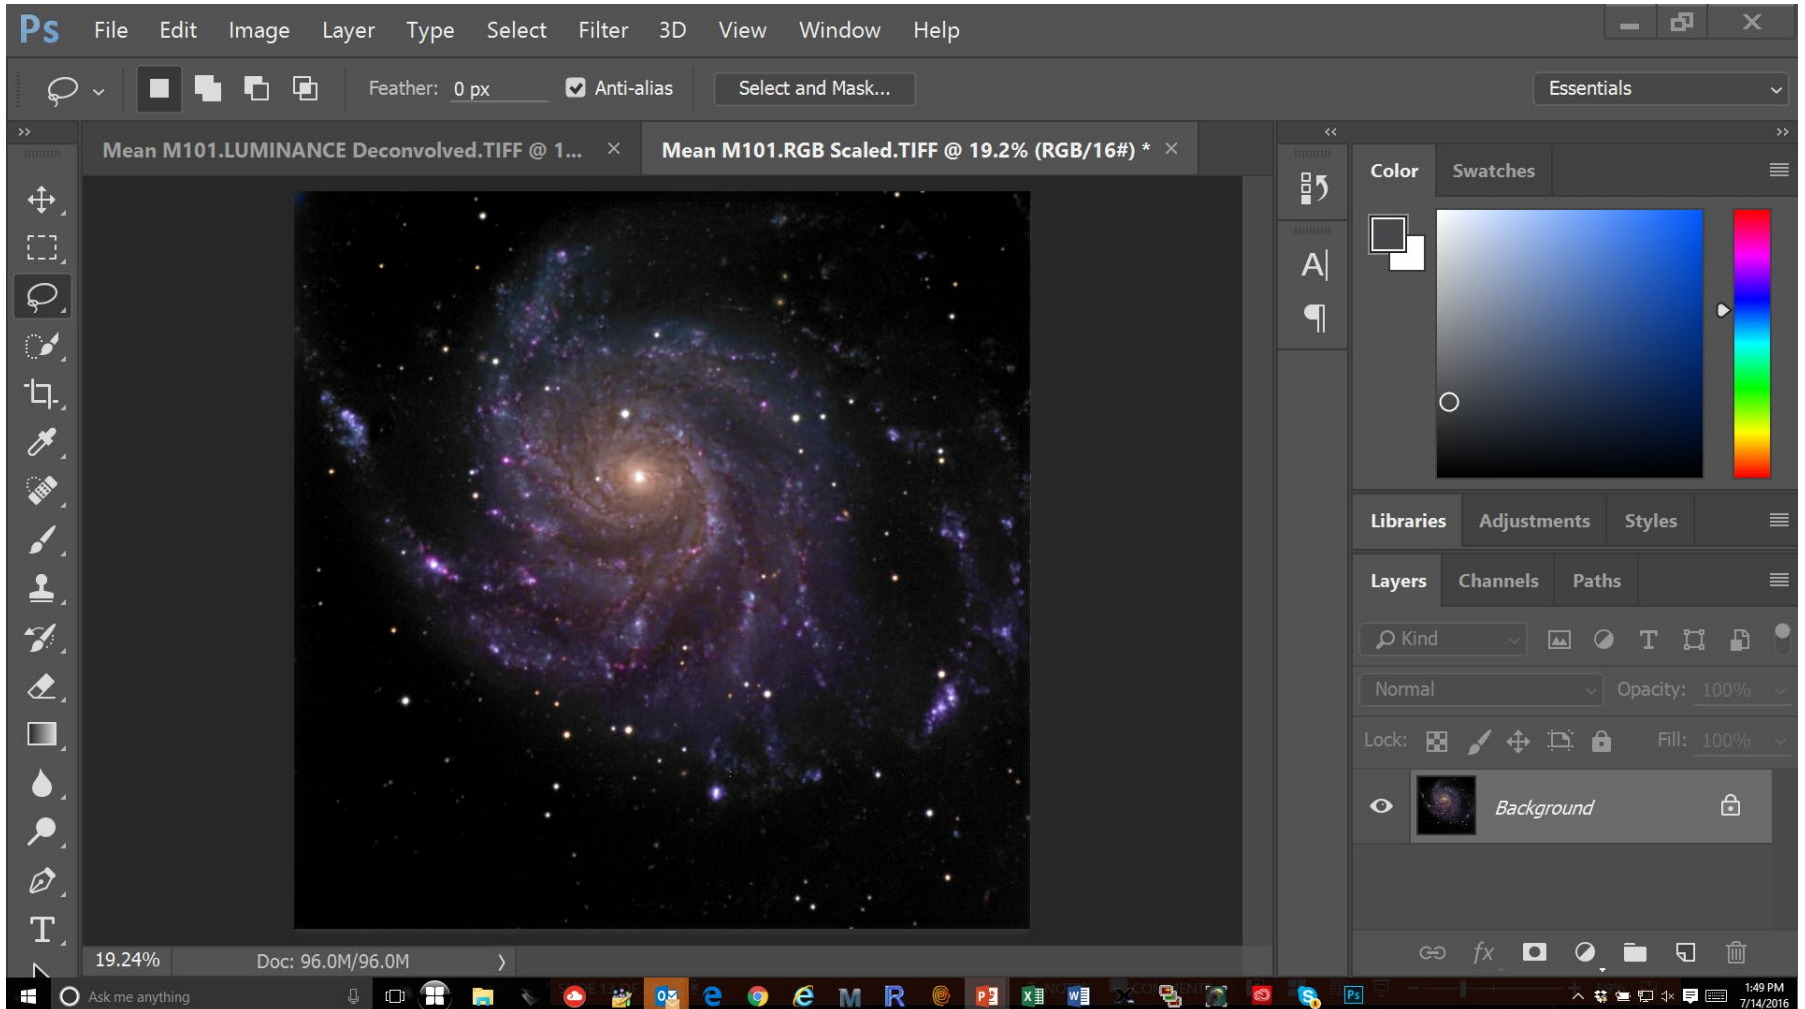


---

# Flattening and retouching


- ❑ When you're happy with the tuned-up image superposition, combine the two layers into one by using [Layer > Flatten Image](#).
  - Now it's back to being a three-plane RGB image but still bears the L image's transmission; hence LRGB.
- ❑ Retouching: necessary to get rid of any remaining flat-field or deconvolution artifacts. This gets into some of the real magic of Photoshop.
  - The [Lasso](#) tool, , and [Edit > Fill > Content-Aware](#). Best thing for dust donuts and other flat-fielding flaws. Enclose them with the lasso, erase them with the menu command.
  - The [Brush](#) tool, , is a versatile way to click away tiny blemishes such as deconvolution artifacts. I still find myself using [Lasso](#) and content-aware fill for this, but the brush is quicker.

# Retouched LRGB image



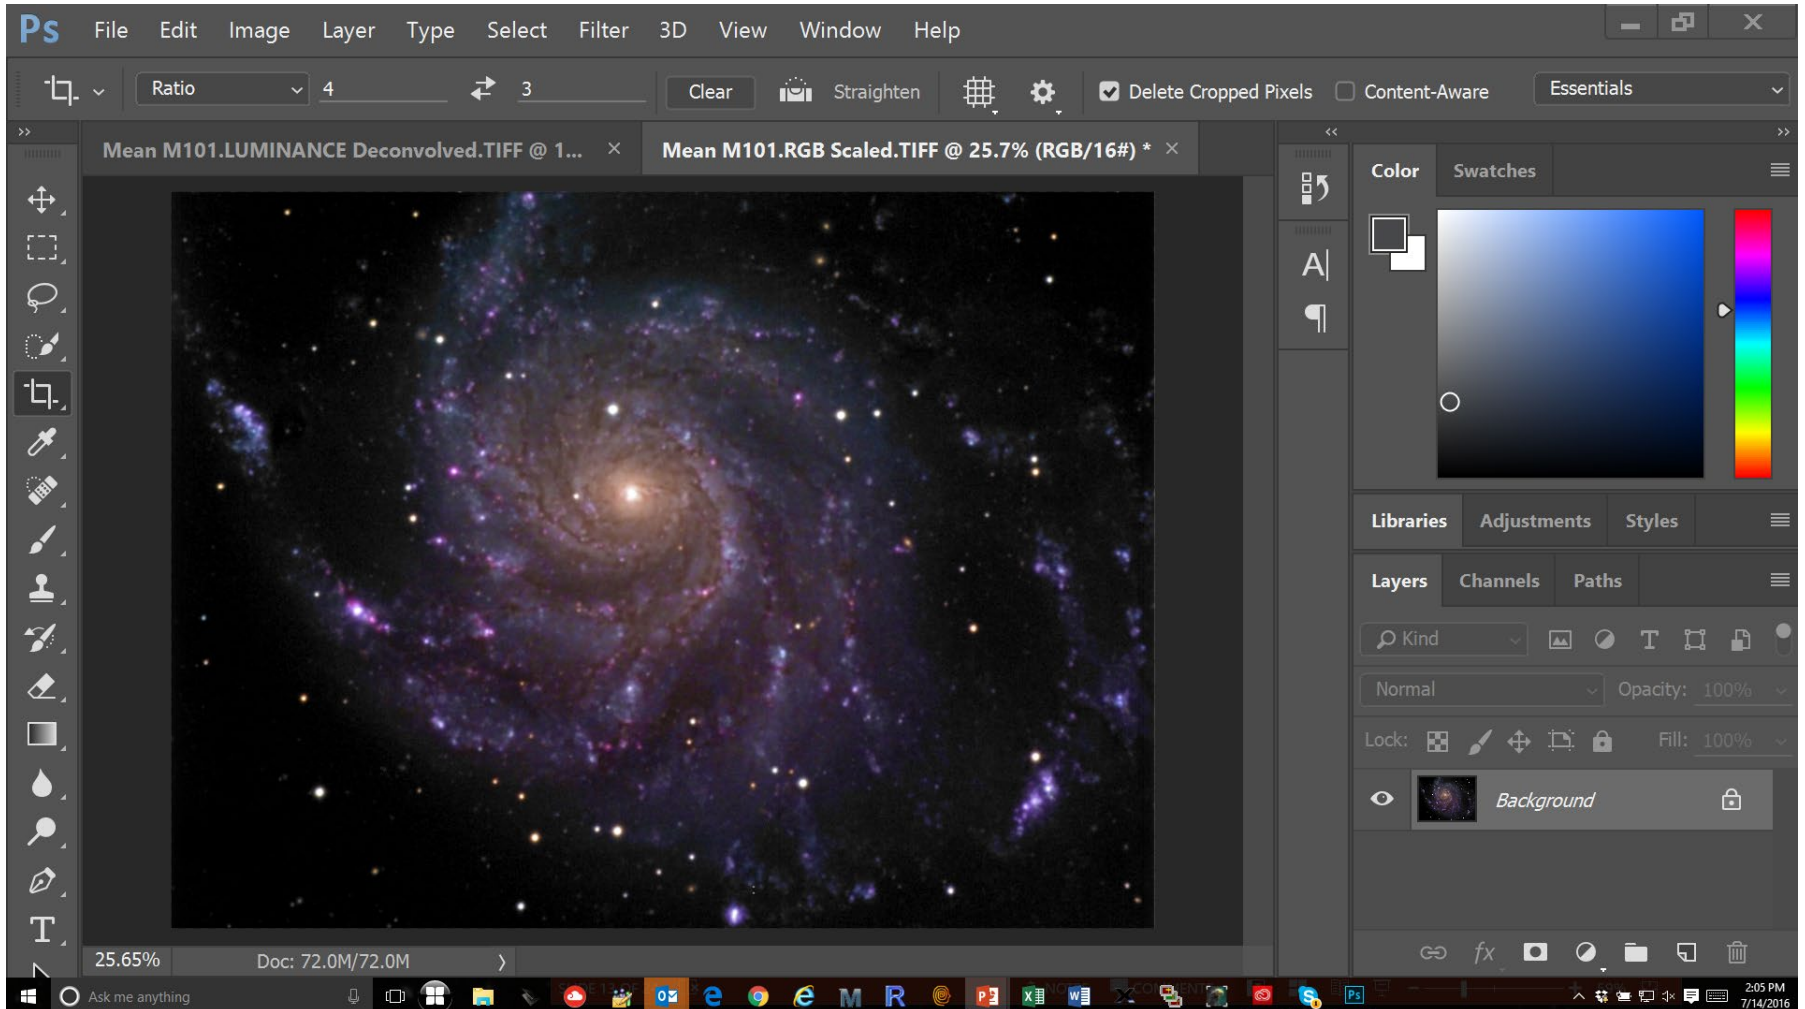
---

## Final steps

- ❑ **Save As...**, of course. Don't overwrite your original TIFF images.
- ❑ Save the full frame before you crop, as a 16-bit TIFF image.
- ❑ Use the **Crop** tool, , to prepare a trimmed version of the image for posting on the Web (PNG) or printing and framing (TIFF).
  - Note that the **Crop** banner has some useful ways to set the cropping frame's aspect ratio precisely.
  - Picture frames frequently have aspect ratios 4:3, 5:4 or square.
  - I use **Fracture** (<https://www.fractureme.com/>), which prints images directly on glass and adds a stiff backing, for nice frameless hangings in many sizes and 4:3 or square aspect ratio.
  - Printing takes practice: looks different on screen and on paper.



# 4:3 cropped image







© 2014 Spring 2010