


# IMAGE CALIBRATION AND ANALYSIS TOOLBOX USER GUIDE



Jolyon Troscianko & Martin Stevens

Contact: [jt@jolyon.co.uk](mailto:jt@jolyon.co.uk)

## Contents

OVERVIEW.....	3
EQUIPMENT.....	5
Objective, camera-specific measurements.....	5
Equipment check-list for objective measurements.....	5
Converting to cone-catch images.....	6
Equipment check-list for UV photography and cone-mapping.....	6
TAKING PHOTOS.....	8
Lighting – emission spectra.....	8
Lighting – direction & diffuseness.....	10
Grey standards.....	11
Camera settings.....	12
Scale bar.....	14
Taking photos check-list.....	15
SOFTWARE.....	17
Software check-list.....	17
Installation.....	17
DCRAW.....	18
Compiling DCRAW for MacOS.....	19
Memory.....	20
IMAGE PROCESSING.....	21
What is a “multispectral image”?.....	21
Checking Photographs – RAWTherapee.....	22
Creating Multispectral Images.....	24
Select regions of interest and a scale bar.....	28
Reloading multispectral images.....	30
Generating a new cone mapping model.....	30
Generating a custom camera and filter arrangement.....	33
IMAGE ANALYSIS.....	34
Batch image analysis.....	35
Batch image analysis results.....	39
Granularity pattern analysis introduction.....	39
Calculating pairwise pattern and luminance differences.....	42
Colour differences introduction.....	43
Calculating pairwise colour and luminance discrimination values.....	43
PRESENTING DATA.....	46
Statistics.....	46
Reporting your methods.....	46
Creating colour or false-colour images for presentation.....	47
REFERENCES.....	50

---

# OVERVIEW

---

The Image Calibration and Analysis Toolbox can be used to transform a normal digital camera into a powerful objective imaging tool.

Uncalibrated digital photographs are non-linear (Stevens et al., 2007), meaning that their pixel values do not scale uniformly with the amount of light measured by the sensor (radiance). Therefore uncalibrated photographs cannot be used to make objective measurements of an object's reflectance, colour or pattern within or between photographs. This toolbox extracts images linearly from RAW photographs (Chakrabarti et al., 2009), then controls for lighting conditions using grey standards. Various colour and pattern analysis tools are included, and images can be converted to "animal vision" images.

As long as none of the photos are over-exposed, the processing will not result in data loss when measuring reflectance values above 100% relative to the standard (which is common with shiny objects, or when the standard is not as well lit as other parts of the image). This is because all images are opened and processed as 32-bit floating point straight from the RAW files as they are required, so no (very) large intermediate TIFF images ever need to be saved.

Using two or more standards (or one standard with black point estimates) overcomes the problem of the unknown black point of the camera sensor, making the method robust in the field with various factors reducing contrast such as sensor noise or lens flare.

Photographing through reflective surfaces, or through slightly opaque media is also made possible by using two or more grey standards. This allows photography through water from the air (as long as the reflected sky/background is uniform), underwater, or through uniformly foggy/misty atmospheric conditions.

Multispectral cameras with almost any number of bands are supported by the code and can be used for greater colour measurement confidence.

This toolbox is freely available on the condition that users cite the corresponding paper published with this toolbox: "Image Calibration and Analysis Toolbox – a free

software suite for measuring reflectance, colour, and pattern objectively and to animal vision", and the papers relating to any visual systems and natural spectrum libraries used.



*Image 1: Examples of flowers photographed in human-visible colours (left) and false-colour honeybee vision (right).*

---

# EQUIPMENT

---

The ideal equipment you need will depend on the hypotheses you want to test.

## Objective, camera-specific measurements

For making objective measurements of reflectance (the amount of light reflected by an object relative to a reflectance standard), colour, and pattern all you need is a digital camera that can produce RAW images, a good quality lens (e.g. with no vignetting – darkening at the image edges), and a grey or white standard (or, ideally two standards – dark grey and white). Most pattern analysis requires a scale bar. Lighting also needs to be considered carefully (see page 8 onwards and Image 3).

The images produced by this set-up will be objective in terms of measuring reflectance levels, being repeatable with the same equipment, and robust against changes in lighting conditions, but are strictly device-specific. Different camera models have different spectral sensitivities, so measurements made with a different set-up will produce slightly different results.

## Equipment check-list for objective measurements

- **Camera body.** Use any camera that can take RAW photos. This includes digital SLRs and mirrorless cameras. Some crossover cameras with integrated lenses might also be suitable.
- **Lens.** Ideally use a “prime” (i.e. not a zoom) lens. A high quality lens is desirable as this will minimise vignetting (images getting darker towards the corners), radial distortion (e.g. if you photograph a chequer board the lines in the photo should all be perfectly straight), and chromatic distortion.
- **Diffuse grey standard(s).** These can be purchased from most photography suppliers and will specify their own reflectance value. The X-rite colorChecker passport is convenient as it has a range of grey levels.
- **Lighting:** Sunlight, camera flash, arc lamps, incandescent bulbs, or white phosphor-based LED lights are suitable. Avoid fluorescent tube lights (see page 8). Use the same light source for all photographs if possible. Try to use a diffuse light source for shiny or 3-dimensionally complex objects (see page 10).

## Converting to cone-catch images

If your hypotheses depend on the appearance of the object to a specific visual system/model species, or absolute measures of colour are required then cone-catch images are recommended.

This conversion produces images based on the spectral sensitivities of a given visual system. The images are device independent (different cameras, or spectrometers should all produce the same results).

The equipment required will depend on the sensitivities of the visual system in question. Many species are sensitive to ultraviolet (UV) light, in which case a UV camera set-up would be required. This entails a camera converted to full-spectrum sensitivity, two filters (passing visible and UV light respectively), and grey standards that have a flat reflectance across the ~300-700nm range (such as sintered PTFE standards). If photographs are being taken in the lab then a UV light source is also required (see below). If your model visual system is only sensitive to wavelengths in the 400-700nm (i.e. human-visible) range, then the equipment listed above for objective images is suitable, with the exception that the spectral sensitivities of the camera must be known.

## Equipment check-list for UV photography and cone-catch mapping

- **Camera converted to full-spectrum sensitivity**. This process involves removal of the filter covering the sensor. Often this filter is replaced with a quartz sheet that transmits light from 300 to 700nm. ACS ([www.advancedcameraservices.co.uk](http://www.advancedcameraservices.co.uk)) and other companies provide this service commercially, or you can try to do it yourself (e.g. [www.jolyon.co.uk/2014/07/full-spectrum-nx1000](http://www.jolyon.co.uk/2014/07/full-spectrum-nx1000)). Doing it yourself could damage the camera, and has the added difficulty of ensuring the sensor position can be adjusted to restore focusing.
- **UV lens**. Most "normal" lenses do not transmit well in UV. CoastalOptics and Nikon make excellent but expensive UV-dedicated lenses. Cheaper alternatives can be purchased second hand online. Well known examples are the Nikkor EL 80mm (older model with a metal body) enlarging lens and the Novoflex Noflexar 35mm. The Nikon AF-S Micro Nikkor 105mm f/2.8 does have UV transmission, although it is not achromatic, meaning it needs re-

focusing between visible and UV photos (which complicates the alignment, although the toolbox can readily deal with this).

- **UV-pass and visible-pass filters.** We recommend the Baader Venus-U filter, which transmits efficiently from ~320 to 380nm. The Baader IR/UV cut filter can then be used for visible-spectrum images (~400 to 700nm). We use custom-built filter sliders to switch easily between visible and UV-pass filters. These sliders are made on a CNC milling machine and G-Code/plans can be made available on request.
- **UV-grade grey standard(s).** Normal photography grey standards are not suitable as few have a flat reflectance down to 300nm. Spectralon standards from Labsphere are suitable. In an emergency or adverse conditions a finely sanded piece of natural white PTFE plastic wrapped in a stretched-flat layer of plumber's PTFE tape could be used as a white standard. Although this will not be quite as diffuse as a Spectralon standard it should have flat reflectance down to 300nm. The layer of PTFE tape can also be replaced regularly as it gets dirty. Never touch the surface of any PTFE/Spectralon grey standards. While they are highly hydrophobic and repel water extremely well they absorb oils incredibly easily (such as the natural oils in your skin). Also take great care not to get sun-block/sun-screen chemicals near the surfaces.
- **UV lighting.** Only use one type of light source that covers the whole spectrum – never attempt to use a light that only emits UV and a separate source for visible light. See below for suggestions (page 8, Image 3). Only use a metal-coated reflective umbrella or sheets of natural white PTFE to diffuse



*Image 2: A UV camera setup with our custom-built filter slider that makes it easy to take the same photo in visible or UV bands.*

UV light sources (avoid using “white” umbrellas or diffusers as they are very unlikely to be white in UV, see page 10). Natural white PTFE sheets of about 0.25 to 1 mm thickness are good for diffusing through back-lighting, 1 mm or thicker sheets are good as reflective white surfaces.

- **Tripod.** It is essential that there is as little movement as possible between the visible and UV photographs so that they can be perfectly aligned.

---


# TAKING PHOTOS

---

## Lighting – emission spectra

Photography is entirely dependant on light and the way it bounces off or through objects, so is an important consideration before starting data collection. In an ideal world, all photographs that you want to compare should be taken under uniform lighting conditions, however in practice this can be impossible to achieve. The emission spectrum of a light source affects its colour, which can generally be controlled for by using a grey standard (see below). However, some emission spectra are so “spikey” that the spikes can interact with reflectance spectra in ways that cannot be controlled for with a grey standard (e.g. causing metamerism, see Image 3). Fluorescent light sources (energy saving bulbs, tube lights etc.) are the worst and should be avoided if at all possible. Flickering can also be problematic. If your photographs have odd looking horizontal banding this is caused by a flickering light source and a much longer exposure (shutter speed) should be used to eliminate this effect if no other source is available. Also be aware that many artificial light sources change their spectra as they warm up (over a few minutes), could change with voltage fluctuations in the mains or battery supply, and can change as the bulb ages. So always try to get a grey standard into each photograph if possible rather than relying on the sequential method (taking a photograph of the standard before or after the target has been photographed).





*Image 3: Emission spectra of different light sources (normalised so that the maximum spectral radiance = 1). The sun is the most ecologically relevant light source and provides a good broad emission spectrum, but this also varies substantially with time of day, latitude, and atmospheric conditions. Arc lamps have relatively "spikey" emission spectra, but this example (Iwasaki eyeColor with its UV filter removed, see below) does a relatively good job of recreating sunlight. Fluorescent tubes are very poor light sources for accurate colour rendering because of their very sharp spikes with some wavelengths almost entirely absent and others over-represented. These spikes can interact with reflectance spectra to make colours look quite different to our eyes or to a camera under sunlight versus fluorescent lighting.*

The light source you use must cover the entire range of wavelengths you are photographing, it is not acceptable to use two or more light sources with different emission spectra to cover the whole range you are interested in unless these light sources can be blended together effectively (which is not trivial). For example, using one light source that emits human-visible light should not be combined with a second light to add the UV component. This is because the lights will interact with the 3D surface angles of the target, making some regions appear to have colours that they do not.

The Iwasaki eyeColour MT70D E27 6500K arc lamp available from CP-lighting ([www.cp-lighting.co.uk](http://www.cp-lighting.co.uk)) can be converted into a good UV-visible band light source by removing its UV/IR protective filter. The filter is just visible as an oily rainbow effect when the bulb is held up to the light. By using a hand-held drill with a steel wire circular brush this filter can be removed without damaging the glass to increase the UVA emissions of this bulb. Sensible precautions and protective equipment (goggles and gloves) should be worn when doing this. The bulb's emissions in the UVA range are increased by this process, so eye and skin protection must be worn when working near this light source for long periods, and it will fade nearby colours faster. Being an arc lamp this bulb needs to be operated by a suitable ballast (also available from CP-lighting). Photography stands that have an E27 socket

can also readily be used with this bulb. Consult an electrician to wire this system together if you are not confident in doing it yourself.

## Lighting – direction & diffuseness


The direction and diffuseness of the light source can interact with the three-dimensional shape of the target, and with the “shinyness” or “glossiness” of its surfaces. A point-source of light (such as the sun, or a small light bulb) is highly directional, so the reflectance measured from a point relative to the grey standard will be highly dependent on its angle relative to the light source as well as the surface reflectance. Therefore, under highly directional lighting the reflectance measured from a surface by your camera will only be accurate if that surface has perfect Lambertian reflectance (is very diffuse), and the angle of the surface is exactly equal to that of the grey standard. Clearly these conditions are almost never true, so care must be taken with directional point light sources. Direct sunlight is arguably a different case because measurements made by a camera in the field will be qualitatively similar to the reflectance measured by an eye, and the lighting conditions are ecologically relevant. However, as long as uniform lighting angles relative to the target are being used, point sources are suitable for comparing reflectance values within a given study.

The effect of the surface angle (i.e. the shape of the target) and its diffuseness can be minimised by using a diffuse light source. Shiny objects and complex 3D objects will therefore benefit most from diffuse lighting conditions. Artificial light sources can be made more diffuse with standard photography umbrellas or diffusers in the human-visible range. However, particular care must be taken when diffusing UV light sources as the diffuser must be UV reflective. White plastic umbrellas and diffusers should therefore not be used for UV sources. Metal-coated umbrellas or sheets of natural white PTFE (Polytetrafluoroethylene, easily bought online from plastic stockists) are suitable for diffusing UV light sources.

In the field, lighting angle and diffuseness can be difficult to control for. If possible, attempt to photograph only under sunny or overcast conditions, not both. If this is not practical, or you are interested in the variation in natural lighting conditions then take note of the lighting for each photograph and take multiple photographs of each sample under all ecologically relevant lighting conditions. Then lighting conditions can be entered into the statistical model during analysis.

## Grey standards

At least one grey standard of known reflectance is required in each photograph (or in a separate photograph taken under exactly the same conditions and with identical camera settings). The position and angle of the grey standards is critically important for objective measurements. The light falling on the standard should match as closely as possible the light falling on the object being measured with respect to angle relative to the light source, intensity, and colour (see page 10, above). If only one standard is being used then a value around 20 to 50% reflectance is appropriate for most natural objects, as a rule the ideal standard should have a reflectance near the upper range of your sample. For example, when measuring very dark objects a lower reflectance standard can be used (e.g. 2 to 10%).



*Image 4: When photographing roughly 2D objects the grey standard(s) should always be in the same plane as that object and ideally the same distance from the light source.*

Multiple standards (e.g. white and dark grey) allow the software to overcome some optical glare (light bleeding onto the sensor), or photographing through reflective or hazy materials such as underwater, and through the surface of water. When photographing from the air into water the grey standards must be positioned next to the target (e.g. both standards underwater as close to the sample as possible), and the background sky or ceiling being reflected by the surface of the water must be uniform (e.g. this won't give correct values if there are any visible ripples or if you can see blue sky and white clouds reflected in the water surface).

When measuring 2D objects the surface of the standard should be in the same plane as the target (e.g. Image 4). For complex objects the standard should normally be angled relative to the light source (e.g. flat on the ground outside, Image 5). Whatever rule is used must be kept constant throughout the experiment and between treatments.



*Image 5: When photographing more complex 3D objects, or when the angles you photograph from cannot be predicted, then the grey standard(s) should be angled relative to the light source. In the field the grey standard should normally be level with the ground so that it collects light from the whole sky, not angled directly towards the sun.*

Remember that light “spills” off nearby objects. Avoid photographing the grey standards near brightly coloured objects that you have introduced to the scene. This includes objects colourful in UV, many surfaces that look white to us absorb UV.

## Camera settings

Focus and exposure are the main considerations when taking photographs. A tripod is recommended if camera shake could affect the photograph, and tripods are essential for UV photography as each photograph needs to be taken twice (once with a visible-pass filter and once with a UV-pass filter) with as little movement as possible between photographs so that they can be aligned.

Exposure is affected by aperture, ISO setting, and shutter speed (also known as integration time). The aperture affects the amount of light allowed through the lens with an iris that changes size like the pupil in your eye. A large aperture (also known as a “wide” aperture, e.g. f/2) lets through lots of light, but that narrows the depth of field, making out-of-focus parts of the photograph even more blurred. Close-up or “macro” photography creates even narrower depths of field, so a smaller aperture (e.g. f/16) might be necessary to get more of the scene in focus.


Aperture can interact with various lens imperfections, such as radial chromatic distortion, so as a rule try to keep the aperture constant across all your photos and rely on shutter-speed to change the exposure. Smaller apertures are recommen-

ded to get more of the scene in focus. Most lenses have an optimal aperture of around f/8 for the sharpest images (assuming the whole image is in focus), smaller apertures can create blurring through diffraction, but will make out of focus objects sharper. ISO should also be kept constant throughout your entire study because this affects the signal-to-noise ratio of the images (higher ISO produces more noise). It is important that the aperture is not changed between visible and UV photos. The exposure can be adjusted by altering the shutter-speed and/or the ISO as long as the same ISO setting is used for all visible photos and another ISO is used for all UV photos.


Over-exposure will result in the brightest pixels in the image being “clipped” or “saturated”. As pixels become brighter they reach this level and cannot go any higher, which is a big problem for objective photography as it will result in false data being produced and any pixels that have reached saturation point (which is 65535 for a 16-bit image) should not be measured. Most digital cameras provide on-screen histograms and these are very useful for judging whether a photo is over-exposed as you take the photos (though they cannot be relied on entirely because they use non-linear levels and are often conservative). Ideally, separate red, green and blue histograms should be used, but some cameras only provide a grey-scale histogram. Single grey-level histograms can be less reliable for UV photography. If the grey values are calculated as the average of RGB this will under-estimate the pixel levels because the green channel is not sensitive to UV, while the red is most sensitive, meaning the red could be saturated. Nevertheless, once you have become accustomed to your camera's histogram performance it is a very useful tool for judging exposure.

The recommended camera mode depends on conditions. In fixed lab conditions where lighting intensity will not fluctuate substantially the best solution is to use “M” (full manual aperture and shutter speed control), and spend some time working out the best settings, then use those settings for the whole experiment. Note that any changes in the position or intensity of the light source will likely affect these settings, so check histograms regularly. Some cameras do a very good job of calculating the optimal exposure, in which case “A” mode (aperture priority) can be used. In this mode you set the aperture to use (which should remain fixed across the whole study), and the camera automatically decides on the best shutter speed. You can then alter the exposure value (“EV” or “+/-” symbol on most cameras), telling the camera to under- or over-expose the photo by a set number of exposure values compared to what the camera thinks is best. For example, if you are photographing something against a dark background the camera will gener-


ally try to over-expose, so selecting EV-1 will tell the camera to under-expose slightly. DSLR cameras tend to be good at working out the optimal exposure in visible wavelengths, but are poor at working out UV exposures, so manual control (M) is recommended for UV photography. Mirrorless cameras (such as the Samsung NX1000) tend to do a better job in UV as their exposures are based on the light hitting the sensor rather than dedicated light meters, so can often be used in aperture priority mode. In either case though it is often possible to use the exposure values as references that will vary between camera models and will require some experimentation to get used to. See Image 6 for examples.




*The camera thinks this is the optimal exposure.*




*The camera thinks this is one stop (EV) over-exposed. This might produce a good exposure for bright targets photographed against a darker background.*



*Exposure bracketing. With this setting three photos will be taken with +/- 1 EV. i.e. one photo under-exposed, one over-exposed and one normal.*



*Exposure bracketing as above, but with a larger EV range. This is recommended in highly unpredictable situations or for very high contrast photos.*



*Exposure bracketing with +/- 1 EV, but this time they are all shifted down one stop. Use this if the camera tends to over-expose.*

*Image 6: Exposure values and bracketing examples. In aperture priority (A) mode you can manually shift the EV values to meet these specifications. In manual (M) model you can change the shutter speed to shift these values.*

Remember that over-exposure of your samples can lead to complete data loss. Automatic exposure bracketing is a useful feature on almost all cameras capable of taking RAW photos whereby the camera automatically adjusts the exposure range across three or more photos to intentionally under- and over-expose by a set number of “exposure values”. Using exposure bracketing is a good idea if there's little cost to taking a couple of extra photos to ensure you get the perfect exposure (you can look at the photos on the computer later to judge which is the best exposure), or when the contrast of the samples being photographed is highly variable. Note that some cameras automatically take all three or more photos with one press of the shutter button, but most for models you need press the shutter

once for each photo.

## Scale bar

Pattern analysis requires all images be scaled to a uniform number of pixels per unit length – a number that will vary with every study (tools are provided to calculate this, see page 37). Photographing different samples from slightly different distances can be controlled for with a scale bar, nevertheless we recommend taking photographs at uniform distances from the subject whenever possible. Particular care should be taken to photograph different treatments at the same distances for unbiased pattern analysis. Position the scale bar level with the sample and photograph from overhead if possible. If photographing from an angle rather than directly overhead a horizontal circular disk could be used as a scale bar (the disk's maximum width will always equal its diameter whatever angle it's viewed from) or place a straight scale bar/ruler side-on to the camera. Even if your initial hypothesis does not concern pattern or size it is always good practice to include a scale bar.

## Taking photos check-list

- Photograph in **RAW format**, not JPG. The camera white-balance is not important (this is ignored during RAW import).
- Use a **fixed aperture** between all photos in the study if possible to have a uniform depth of field – use changes in shutter speed to control the exposure.
- When using a zoom lens only use the maximum or minimum focal length, not an intermediate.
- Ensure the **lighting** is suitable and the angle of the **grey standard(s)** relative to the light source follows a clear, justifiable rule (see Image 13), and that the colour and brightness of the light falling on the standard is as close as possible to that of the sample.
- For pattern analysis or scale measurements place a **scale bar** level with the sample and photograph from a consistent angle (e.g. overhead).
- Place a **label** in each photograph if practical and make a note of the photo number linked to each sample.
- If you are placing samples against an **artificial background** use a dark grey or black material – particularly with potentially transparent or “shiny” samples. When photographing samples *in-situ* in the field their natural background is normally most suitable unless you are comparing the sample

directly to its surroundings, in which case it may be desirable to photograph the sample separately against a dark background so that it is measured independently of its surrounds (obviously only if it can be moved without risk of harm to yourself or the sample).

- Use a stable **tripod** & remote shutter release if possible (tripods are essential for UV photography).
- Consider using **exposure bracketing** to get the best exposure and regularly check photo histograms on the camera to make sure they are well exposed. The bars of the histogram should be as evenly spread from left to right as possible without quite touching the right-hand side.
- Always check the photos “look good”, that everything important is in focus, and fills the frame as much as possible.



---

# SOFTWARE

---

## Software check-list

- **ImageJ** version 1.49 or more recent with Java, available from [imagej.nih.gov/ij/](http://imagej.nih.gov/ij/) (Schneider et al., 2012). 32-bit or 64-bit versions are supported. Ubuntu: available from the software centre. Note that on Linux ImageJ is more stable with Java version 6, though it does generally work with version 7.
- Download the **Toolbox** files from [www.sensoryecology.com](http://www.sensoryecology.com) or [www.jolyon.co.uk](http://www.jolyon.co.uk). Unzip the contents of the zip files and place them in your imagej/plugins directory (see below for further details).
- (Optional) **R**, from [www.r-project.org/](http://www.r-project.org/) (R Core Team, 2013), required for creating new cone-catch mapping models for converting from camera-vision to animal-vision. Once the models have been made for a given camera and visual systems R is not required.
- (Optional) **RAWTherapee** from [rawtherapee.com](http://rawtherapee.com). This is useful for screening RAW photographs before processing.

## Installation

Install ImageJ, following the instructions for your operating system. Download the toolbox files from [www.sensoryecology.com](http://www.sensoryecology.com) or [www.jolyon.co.uk](http://www.jolyon.co.uk), unzip the contents and copy them to your ImageJ plugins directory.

On windows this directory is normally **C:/Program Files/ImageJ/plugins** (see Image 7). On MacOS this is **Applications/ImageJ/plugins**. On Ubuntu this is **home/username/.imagej/plugins** if you have installed ImageJ through the software centre (note this is a hidden folder on Ubuntu).




Image 7: Extract the toolbox files from their zip file and copy them to your imagej plugins folder

## DCRAW

Each camera manufacturer has their own type of RAW file and DCRAW is an incredibly useful piece of software that can open almost all types of RAW file (Dave Coffin [www.cybercom.net/~dcoffin/dcraw/](http://www.cybercom.net/~dcoffin/dcraw/)). Our toolbox utilises DCRAW to open RAW files in a linear fashion (Chakrabarti et al., 2009), using the DCRAW Plugin for ImageJ (Jarek Sacha, [ij-plugins.sourceforge.net/plugins/dcraw/](http://ij-plugins.sourceforge.net/plugins/dcraw/)). Incidentally, RAWTherapee also uses DCRAW to open RAW files, meaning it is suitable for pre-screening photographs to check whether they are well exposed (see below).

To check DCRAW is working on your installation of ImageJ go **Plugins>Input-Output>DCRAW Reader**. Select a RAW file (e.g. .NEF for Nikon, .CR2 for Canon), and it will show an options dialogue. Select the options shown in Image 8 and it should load the RAW file as a linear 16-bits per channel image. Note that ImageJ treats 16-bit and 32-bit RGB images as a stack displayed in colour. This is slightly conceptually different to a normal 24-bit RGB image (8-bits per channel), and any measurements you make of this image will only be from the selected channel (even though you can see the channels).




Image 8: DCRAW Import options. The settings specified here will open a linear image from the selected RAW file. If the toolbox is not working correctly try opening a file with DCRAW to see if this is the problem.

## Compiling DCRAW for MacOS

DCRAW for Windows and Linux seems to work reliably across different operating systems, however, MacOS sometimes requires recompilation of DCRAW. If the Toolbox does not work with the supplied files and throws out an error message when opening a RAW image you will need to compile your own DCRAW binary file (this is the bit of software that converts the RAW image into a usable image for ImageJ).

1. Create a new folder on your desktop (e.g. called myDCRAW).
2. Get the original source code for DCRAW from this website: <http://www.cybercom.net/~dcoffin/dcraw/dcraw.c>, in your browser go File>Save Page As.. and save the code as "**dcraw.c**" in your "**myDCRAW**" folder.
3. Open **System Preferences**, Choose the **Keyboard** option and then the "**Shortcuts**" tab. Under "Services", check the **New Terminal at Folder** option. Close System preferences.
4. **Right-click** your **myDCRAW folder** and choose the "**New Terminal at Folder**" option. A terminal will load up, you can check it's in the right location by typing "ls" and hitting enter, it should list dcraw.c in response.
5. Paste the following command into the terminal and hit enter:

```
llvm-gcc -o dcraw dcraw.c -lm -DNO_JPEG -DNO_LCMS -DNO_JASPER
```

6. As you hit Enter, the computer will download Xcode.app, compile and install the source code. Ignore the warning messages on the terminal. When it is done, check the myDCRAW folder and there should be a file (grey with green "exe" label) named 'dcraw'.
7. Copy this **dcraw** file (not dcraw.c) across to your **imagej/plugins/dcraw** folder, overwriting the existing file.
8. Check whether it works in ImageJ by going **Plugins>Input-Output>DCRAW Reader** and try to load a RAW image with the settings shown in Image 8.
9. If it doesn't work it's possibly because the executable dcraw doesn't have permission to run. To give it permission enter this into the terminal (you may need to specify a different path if you've placed imagej elsewhere):

```
chmod 777 /Applications/ImageJ/plugins/dcraw/dcraw
```

## Memory

Some multispectral images can take up a lot of memory. If you get an error message complaining that ImageJ is out of memory you can increase the memory dedicated to ImageJ in **Edit>Options>Memory & Threads**. Depending on how much memory is available on your computer you can increase the number (e.g. 3000MB should be plenty for most multispectral images).


---

# IMAGE PROCESSING

---

## What is a "multispectral image"?

1. RAW photographs extracted as linear 16-bit images



2. The grey standard location(s) are selected manually

A close-up of a flower petal with a dashed circle around a grey standard. A mouse cursor is pointing at the standard.

3. Photograph channels are aligned automatically or manually. The grey standards are measured following alignment

Two images are shown being aligned. Arrows indicate the movement of one image relative to the other. A grey standard is visible in both images.

4. A ".mspec" file is created alongside the RAW photos, saving the settings required to recreate the normalised, aligned multispectral stack

A screenshot of a window titled 'Image Directory' showing a list of files: visPhoto.RAW, uvPhoto.RAW, and myImage.mspec.

5. The multispectral image stack is converted to 32-bit floating-point and normalised, ready for objective measurement or conversion to cone-catch quanta

A stack of three images (red, green, blue) is shown, representing the multispectral image stack.

6. Regions of interest and a scale bar can be selected. These ROIs are saved with the .mspec image for automatic measurement. Regions with the same prefix can optionally be measured as one

A flower image with three regions of interest labeled p1, p2, and p3, and a scale bar labeled c1.

7. The Batch processing tool can be used to measure multiple images using any selected visual system

*Image 9: Overview of the image processing work-flow.*

A multispectral image is a stack of images captured at different wavelengths. A normal RGB photo is a multispectral image with three wavelength bands captured (red, green, and blue). This software can handle a large number of additional image channels with different wavelength bands represented. The most common extension is into the UV, so in this case we give each channel a name with a prefix that specifies the filter type used (v for visible, u for UV), and a suffix that describes the camera's channel (R for red, G for green, B for blue). Most UV cameras are sensitive to UV in their red and blue channels and not their green channels. So the green channel is thrown out, leaving us with vR, vG, vB, uB, uR. Here the channels are arranged from longest wavelength to shortest wavelength (see Image 9).

Because UV multispectral images are combinations of two photographs, any slight camera shake or change in focus can cause misalignment between the photos. Even an offset of a few pixels can result in false colours being created. Think of the 'purple fringing' seen at the corners of many photos where there is very high contrast, this bright purple is an unwanted artefact caused by chromatic radial distortion, effectively a misalignment between the red, green and blue channels.

## Checking Photographs – RAWTherapee

Before converting and calibrating photos it is often necessary to select the photos with the best exposures (e.g. if exposure bracketing was used), or to check that they are in focus etc. A good bit of software for doing this is RAWTherapee (open source and available across all operating systems).

After installing and opening RAWTherapee ensure the processing profile is set to “Neutral”. We would also recommend setting this as the default profile in **Preferences>Image Processing tab>** under “**Default processing profile for RAW photos**” profile select **Neutral**. Using any other profile might switch on automatic exposure compensation, which makes it more difficult to work out whether a photo is over-exposed.




Image 10: RAWTherapee can be used for screening RAW photos. Ensure the processing profile is set to "Neutral" with no exposure compensation applied.

The RGB histograms in RAWTherapee can be used to determine which photos

have the optimal exposure. These histograms show the number of pixels in the photo across a range of intensity levels from zero on the left to maximum (saturation) on the right. The ideal exposure for the whole image will have a histogram with peaks spread evenly from left to right without any pixels quite reaching the right-hand side (implying they are saturated). If the sample you are measuring is not as bright as other parts of the image then you can check whether it appears over-exposed by running your mouse over the region you will be measuring. As you do this the RGB levels under the cursor will be shown below the histogram. If these values reach the right-hand side of the histogram they are saturated and this region should not be measured.

Image 11 (right) shows a histogram of a well exposed photograph. The peaks on the right of the histogram correspond to the pixels of the white standard, while the large number on the left represent the dark grey background used in this photograph. Image 12 (below) is an over-exposed example of the same image with clustering on the right-hand side of the histogram.

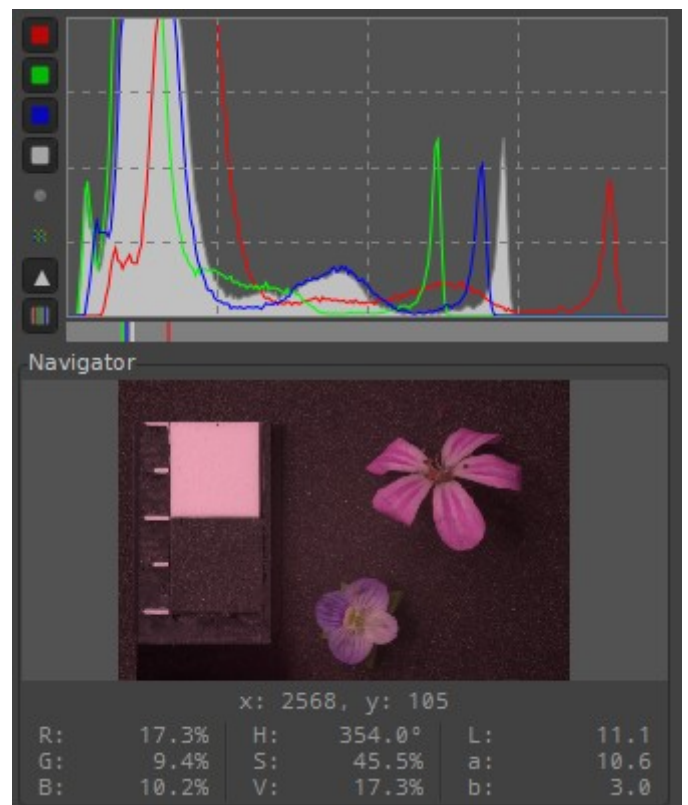


Image 11: Histograms typical of a good exposure. There are no pixels clustered on the very right-hand side of the histogram, but the red peak is close.

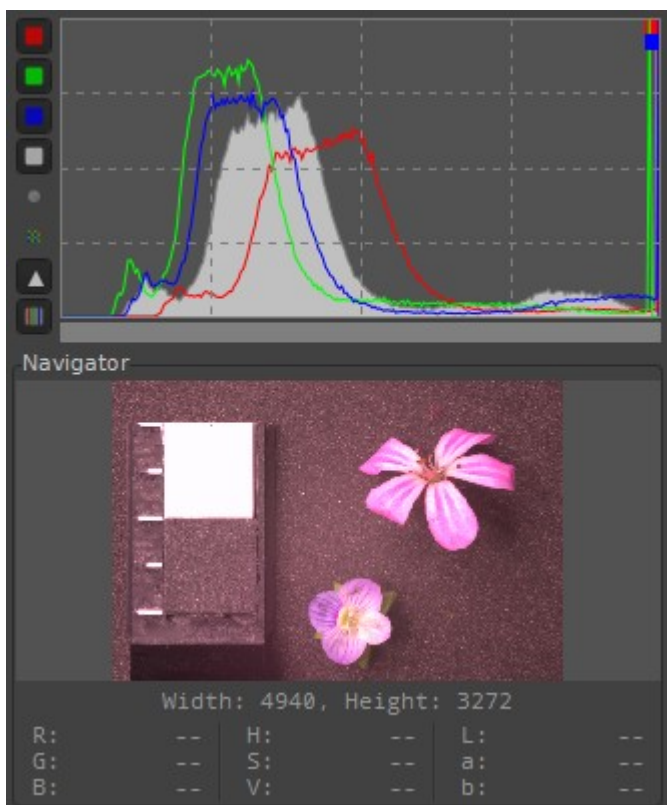
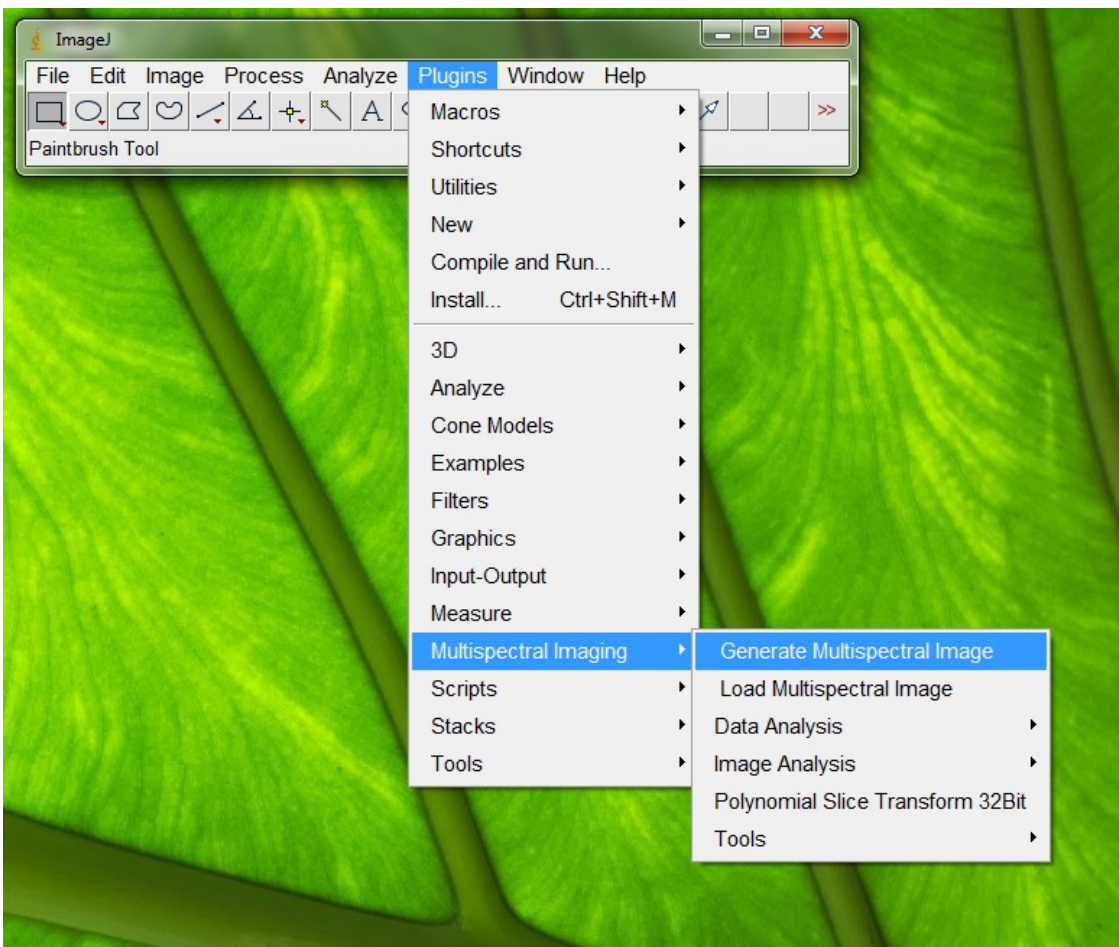


Image 12: Histograms of an over-exposed image, with lots of pixels clustered on the very right-hand side of the histogram.

rocessing

## Creating Multispectral Images

Once you've selected the photos with the best exposure you are ready to generate a multispectral image. For human-visible spectrum photos (i.e. "normal" RGB photos) this will be a single photo, for UV photography this will generally be two photos – one taken through the visible-pass filter and a second through the UV-pass filter. Other filter combinations are supported.



*Image 13: Creating a new multispectral image.*

In ImageJ go **Plugins>Multispectral Imaging>Generate Multispectral Image** (Image 13) this will open a settings dialogue box.

**Settings:** These options specify what type of multispectral stack to create and how the channels should be arranged. The defaults are "Visible" for working with normal RGB photos, or "Visible & UV" for standard UV photography (this will arrange the channels as vR, vG, vB, uB, uR). See page 34 for details on creating a customised filter combination.



**Grey Standards:** Choose “Separate photos” if you did not take photos of the grey standards in the same photo as the sample. This is the sequential method and the photos must be taken with exactly the same camera settings and lighting conditions (shutterspeed, aperture, ISO etc...)

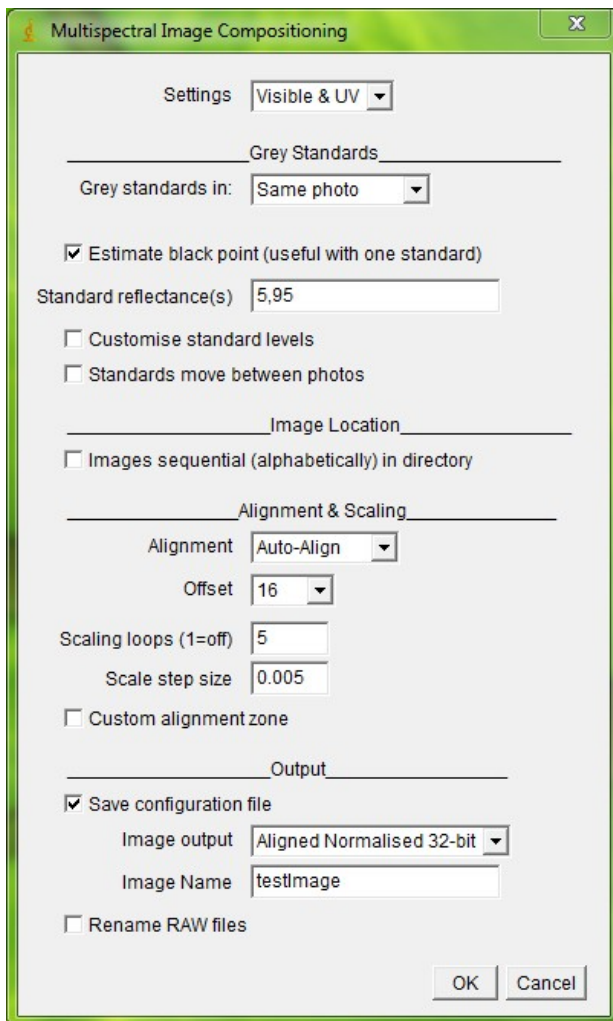


Image 14: Multispectral image settings options.

used for comparing lighting between photographs.

**Customise standard levels:** Grey standards should ideally be grey, meaning they have uniform reflectance across the whole spectrum being photographed (e.g. the reflectance from 300nm to 700nm is uniform for UV photography). However, sometimes standards can get damaged, or maybe two standards were used and they have very slightly different colours. Use this option to correct for known differences in standard reflectance in each of the camera's channels. You can measure these values by photographing your imperfect standard against a reference (clean) grey standard and normalise the image using the clean reference stand-

**Estimate back point:** This setting is useful to minimise optical veiling glare where internal reflections within the lens reduce a photo's contrast, raising the black point from zero. If you only used one standard check the “Estimate black point” box and it will attempt to calculate the dark level this is most important in bright field conditions. *Do **not** tick this box if you have used two or more standards.* It is probably not required when photographs are taken in dark rooms.

**Standard reflectance(s):** Add your standard reflectance values separated by commas if there's more than one. Any number can be used in any order. This example (Image 14) shows the values for a photo that contains a 5% and 95% standard. Leaving this box empty will mean that no standards are used and the resulting image will be linear but not normalised. Linear images such as this can be

ard to work out the corrected reflectance values for each channel. For example, if you measure a damaged 40% standard it might only have 35% reflectance in the camera's uR channel relative to the reference standard, 37% in the uB channel, 39% in the vB channel etc. Measure each channel and input these values when prompted. This setting will rarely be required and is included primarily for emergency use to recover values if a standard is damaged.

**Standards move between photos:** Tick if the standard was moved between visible and UV photos. Otherwise the UV standard locations are measured after alignment to make sure exactly the same areas are being measured between different filters. This setting will rarely be required.

**Images sequential in directory:** If you have arranged your visible and UV photos in the same directory so that the UV photo always follows the visible photo alphabetically this box can be ticked and it will select the UV photo automatically. This also applies for more complex filter combinations and can save a bit of time searching through folders for the correct photos.

**Alignment & Scaling:** These options are ignored for visible-only photography – they are used with UV photography or any other number of filters to align the photographs.

**Alignment:** There is almost always some misalignment between visible and UV photos which needs correcting. The “manual alignment” option will let you manually align the photographs by eye by dragging the image with your mouse. Click on the plus and minus symbols to change the scale to get a match. The easiest way to use this is to align the top left of the image, then change the scale value until the bottom right is aligned too (see Image 15).



Click & drag to align the layers so there is no blue or yellow fringe. Align the top left of the image first.

Change the scaling of the layers using the plus and minus buttons until the whole image is aligned. Repeatedly align the top-left and scale until the bottom right is also aligned.

When the alignment looks perfect (no blue or yellow fringing anywhere) click "Accept".

Image 15: Example of the manual alignment process.

“Auto-alignment” attempts to automatically find the best alignment between photos, but can take a minute or two depending on the photo resolution and processor power. Automatic alignment works best on photos or regions of photos that have lots of detail shared between the visible and UV (or other filter) photos, and does not work so well if your images have large plain areas. Image 16 shows the output of auto-scaling, with the best scale and x-y offsets.

**Offset:** This is only used if “Auto-alignment” was selected, and specifies how far out of alignment the photos can be. 16 pixels is a sensible default, warning messages will be shown in the auto-scaling log (Image 16) if this offset doesn't seem to be sufficient.

**Scaling Loops:** If you have had to re-focus between visible and UV shots this has the effect of zooming the image slightly, which needs to be undone for a true alignment. This is most important when using lenses that are not UV achromatic (like the Nikon Nikkor 105mm, but not the CoastalOpt 60mm). 5-6 scaling loops are normally sufficient. Turn off scaling by setting this value to 1.

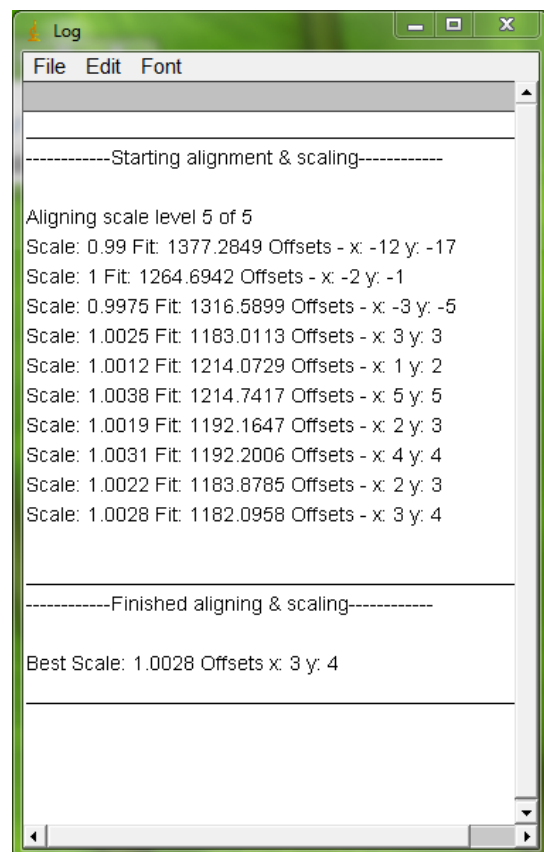


Image 16: Auto-scaling output.

**Scale step size:** A default of 0.005 is sensible (this is a 0.5% scale difference). In searching for the optimal scale this step size is used to start with, and is then halved with each scaling loop.

**Custom alignment zone:** This lets you select the area to use for alignment (for both auto or manual alignment). Select an area with lots of in-focus detail and/or your sample. The grey standards are also useful to use as custom alignment zones. This also speeds things up, because otherwise the entire image is used for alignment. Using scaling is not recommended if you are selecting a relatively small custom alignment zone.

**Save configuration file:** Always leave this ticked unless you are just testing. This specifies whether the .mspec file should be saved. The .mspec file is what's used to reload a multispectral image from RAW files and is saved alongside the RAW files.

**Image output:** Specify what to output at the end for inspection. For UV photos the "aligned normalised 32-bit" option will let you inspect the alignment at the end. "psuedo uv" will show you the G, B and UV channels as RGB (looks good and lets you "see" UV colour, ignoring red). The colour (visible or psuedo uv) outputs can make life easier when selecting areas of interest. *Any output other than aligned normalised 32-bit images are only for inspecting and selecting regions of interest.* Any modifications to these images (e.g. increasing the brightness to see dark objects better) will not be saved and will not affect the multispectral image if it is reloaded later.

**Image name:** Give each multispectral image a name that describes its sample and treatment. This name will become the label used in batch processing. You'll be asked before overwriting in case you forget to change the name at this stage.

**Rename RAW files:** Choose whether to rename your RAW photos with the image name (selected above). Renaming these files can be useful to show you which RAW files have been processed and are associated with which .mspec file. Only use this option for renaming RAW files, if you rename the files yourself the .mspec file will not be able to find them. Always keep the .mspec files in the same folder as its associated RAW photos.

**Click OK** when you are done. The settings you specify here will be saved as de-

faults for next time you make another image.

The script will guide you through the process of selecting the RAW image(s), highlighting the grey standards (you can use a rectangle tool, circle tool, or polygon tool), and optionally

**ALWAYS CHECK THE ALIGNMENT FOR UV PHOTOGRAPHY:** Once the multispectral image has finished it will show you the result. If you selected “normalised 32-bit” as the output format then the script will automatically flip between the various channels to and you can see how well aligned they are (you can also manually flip between the channels using your scroll wheel or keyboard left & right arrows). If there is movement between these images then the alignment has not worked and needs to be repeated (e.g. using manual alignment). If you selected “pseudo UV” as the output then misalignment will look like the blue channel being out of alignment with the red and green.

### Select regions of interest and a scale bar

Regions of interest (ROIs) are used to specify regions of the image that you want to measure. Once you have created a multispectral image you can specify these regions using any of the selection tools (rectangle, circle, ellipse, polygon, freehand tools etc...) After highlighting a region press a key on your keyboard to assign that ROI to a letter. When you are done press “0” to save the ROIs with this multispectral image. For example Image 17 shows an image of two flowers. Using the polygon tool one petal can be selected and then the letter “F” on the keyboard will call that ROI “f1”, highlight the next petal, press “F” and it will be labelled “f2” etc... For the second flower the petals are marked with “G” instead. This will make it easy in the analysis later on to compare the pooled F regions compared to the pooled G regions (flower F compared to flower G), but will also still allow easy between petal comparisons, e.g. if you wanted to look at within-flower differences.

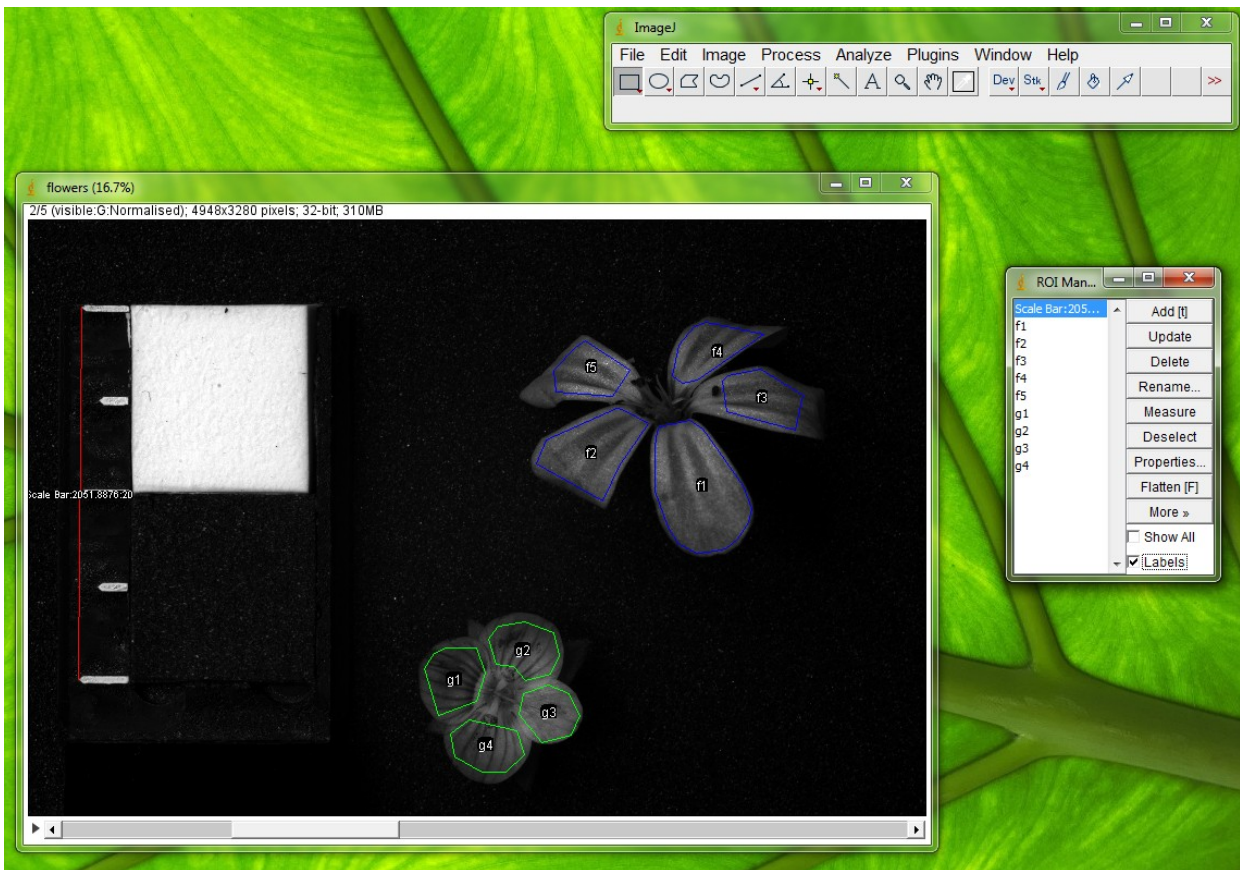


Image 17: Selecting ROIs and a scale bar.

If you are planning to compare patterns between multispectral images, or want to scale your images to a uniform number of pixels per unit length then you need to add a scale bar to your image. To do this select the line tool and draw a line along the scale bar in your image and press "S". A window will pop up asking you how long the length you selected was. Type in a number only (no text). The units of the number used here must be the same for the whole study. E.g. the scale numbers across the whole study should all be in mm, cm or m, but not a combination.

For measuring eggs the toolbox includes an egg shape and size calculation tool that also makes it easy to select an egg-shaped ROI (Troscianko, 2014). Select the multipoint tool (right-click on the point tool and select multipoint if it is not already selected), then place points on the tip and base of the egg and three more points down each side of the egg (8 points in total, though you can add more points). Then press "E" and it will highlight the egg shape. If you are happy with the fit of the egg shape click "accept" and the script will add this ROI in addition to calculating various egg metrics such as length, width, volume surface area and pointedness (shape).

Remember to press "0" when you're done selecting ROIs and the scale bar. If you

leave the log window open it will automatically know where the relevant .mspec file is, otherwise you will be asked which .mspec file to add these ROIs to.

## Reloading multispectral images

Once you've generated a multispectral image you can easily re-load it by going **Plugins>Multispectral Imaging>Load Multispectral Image**.

Select the .mspec file you generated and it will automatically load the normalised, aligned 32-bit image with all the ROIs you selected. You can then use **Plugins>Multispectral Imaging>Tools>Convert to Cone Catch** to manually transform individual images from camera-vision to animal-vision if you have the appropriate mapping function.

## Generating a new cone mapping model

We have included a number of models for converting to animal cone-catch values based on the camera setups that we have calibrated, and under D65 lighting conditions. However, tools are included for making your own model if you know your camera's spectral sensitivity functions.

Mapping from camera to animal vision is performed by simulating the animal's predicted photoreceptor responses to a set of thousands of natural spectra, and then the camera's responses to the same spectra – all under a given illuminant. A polynomial model is then generated that can predict animal photoreceptor cone catch values from camera photoreceptor values (Hong et al., 2001; Lovell et al., 2005; Stevens et al., 2007; Westland et al., 2004). ImageJ gets R to perform the modelling after preparing the data, so you must have R installed for this to function.

You only need to create a mapping function for each camera/visual system/illuminant/spectrum database combination once. Run **Plugins>Cone Mapping>Generate Cone Mapping Model** to create a new model and specify your settings (Error: Reference source not found).

wavelength	300	301	302	303	304	305	306
vR	0	0	0	0	0	0	0
vG	0	0	0	0	0	0	0
vB	0	0	0	0	0	0	0
uB	0	0	0	0	0	0	0
uR	0	0	0	0	0	0	0

Image 18: Camera sensitivity example data.

**Camera:** Select your camera configuration – this will be specific to your camera, lens, and filter combination. You can add your own camera sensitivity curves, visual system receptor sensitivities, illuminant spectra or spectrum database to the relevant folder in `imagej/plugins/Cone Mapping/`. The format of these files should match the

supplied .csv files and the wavelength range and intervals must all match (e.g. 300-700nm or 400-700nm at 1nm increments). The Camera sensitivity .csv file should contain a row for each of the channels required. The order of these channels must match the output of the multispectral image creation. We have ordered channels

from long to short wave sensitivity, e.g. R,G,B, or vR, vG, vB, uB, uR. Where the lower case v and u refer to visible and UV pass filters and RGB refers to the camera's red green and blue channels.

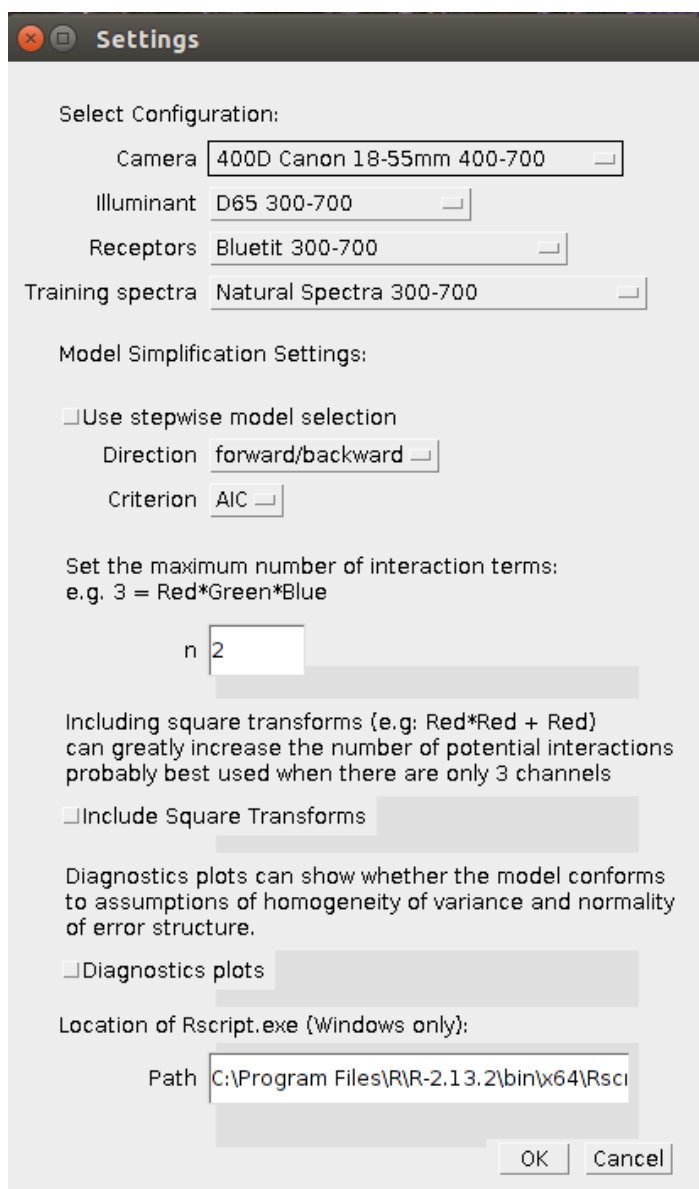


Image 19: Options dialogue for generating a new cone-catch mapping model.

The format for all sensitivity data follow the structure on the left (Image 18), saved as a .CSV file. The order of the camera channel sensitivities is important, and must match the order used when generating the .mspec file (see Error: Reference source not found). This example shows the order for a standard visible & UV multispectral image, starting from the longest wavelength and going to the shortest. See the supplied examples in `imagej/plugins/Cone Mapping`. All sensitivity data (photoreceptors and camera photosensors) must be normalised so that the sum for each row equals 1 (e.g. create these values by dividing each cell by the sum of all the cells across all wavelengths for that receptor). Natural spectra should be normalised so that the highest value equals 1 (e.g. divide



each cell by the maximum value across all wavelengths for that spectrum).

**Illuminant:** Select your illuminant as above. D65 is recommended as the CIE standard illuminant. You can add your own illuminant spectra (see above).

**Receptors:** Specify the visual system to map to following the format in Image 20. This example shows bluetit sensitivities. The order of the channels is not important for photoreceptor classes, but the order here will be other order in which channels are created when converting to cone-catch images.

wavelength	300	301	302	303	304
uv	0.000027	0.0000335	0.0000432	0.0000582	0.0000818
sw	0	0	0	0	0
mw	0	0	0	0	0
lw	0	0	0	0	0
dbl	0	0	0	0	0

Image 20: Sample format for sensitivity data.

**Training spectra:** Select the database of natural spectra to use when making the mapping function. The wavelength band (e.g. 400-700nm or 300-700nm) should be the same for camera sensitivities, illuminant spectra, receptor sensitivities and training spectra. There will be a warning if the wavelengths do not match.

**Use stepwise model selection:** Model selection can be used to remove terms from the model that do not usefully improve the fit. This will slow down the processing of generating the model, but results in a smaller (faster) model for image processing. In practice it seems to make little difference. You can specify how the models are simplified (e.g. forwards/backwards), and whether to use AIC or BIC.

**Number of interaction terms:** The polynomial can specify higher level interactions between channels, though in practice two or three are sufficient.

**Include square transforms:** Including square transforms creates a more complex model, in theory it can provide for a better fit, but in practice it rarely improves the model and can create more noise.

**Diagnostic plots:** These show how well the model fits assumptions of parametric models and are output to imagej/plugins/Cone Models/Plots.

**Rscript Path:** If you are running in Windows you need to add the location of your Rscript.exe file that's bundled with R and performs script processing. If the model

creation process stops doing anything when it says “Waiting for R to process data” there is most likely a problem communicating with rscript.exe – check this file path is correct.

After creating a model it is saved in imagej/plugins/Cone Models/ along with its source code and the quality of the model's fit ( $R^2$  values) are shown for each photoreceptor. The  $R^2$  values should ideally be  $\geq 0.99$ . Considerably lower values imply the model has a poor fit and might fail to produce reliable results.

## Generating a custom camera and filter arrangement

Photo	R	G	B	align1	align2
visible	1	2	3	0	0
uv	5	0	4	3	1

*Image 21: Filter arrangement example for Visible & UV photos.*

You can specify your own channel import options (e.g. if you are using different filter combinations or want to drop the UV blue channel for example). The files specifying this are in imagej/plugins/Multispectral

Imaging/cameras. Add your own .txt file with the following format to specify your own camera/filter setup (Image 21). Each row specifies a filter type, numbers in the RGB columns specify the position that channel should occupy in the stack (zero ignores the channel), and the alignment channels specify which two channels to use when aligning the photos from different filters (the first row is the reference image, so both values are zero). Align1 is the global image number from 1 to 5 in this case, align2 is the channel number relative to this filter (i.e. red=1, green=2, blue=3).

To remove the uB channel in this example, you would change uv-B from 4 to 0, and uv-R from 5 to 4. In this case alignment is performed between channel 3 (visible-B) and uv channel 1 (uv-R), so this can be left as it is.

---

# IMAGE ANALYSIS

---

After generating a multispectral image and opening it as a 32-bit normalised image you can measure the pixel values manually and these will be objective reflectance values, though the colours will be camera-specific. By default the images are in the 16-bit range, from zero to 65535, although being floating point the numbers can go higher if there are parts of the photo with higher reflectance than a white standard. You can divide pixel values by 655.35 to give percentage reflectance values relative to your grey standard(s). By default the channels are shown individually as greyscale images. To make an image with three channels colourful you can go **Image>Colour>Make Composite** and select colour. This will not affect image measurements.

You can manually convert an opened .mspec image to animal cone-catch quanta with **Plugins>Multispectral Imaging>Image Analysis>Convert to cone catch**, and measure the values in these images. Running **Plugins>Measure>Measure All Slices** once after loading ImageJ will measure the pixels in each image channel separately, and now when you press "M" on the keyboard it will measure all the channels. However, we would recommend performing a batch image processing job once you have finished preparing all your .mspec images. This will ensure each image is measured with exactly the same settings and will arrange the data into easily managed spreadsheets with less room for human error.

## Batch image analysis




Image 22: Batch image analysis.

To measure the colour, pattern and luminance data in a series of multispectral images use **Plugins>Multispectral Imaging>Image Analysis>Batch Multispectral Image Analysis** (Image 22). Make sure you select a folder that contains all your .mspec files alongside their associated RAW files.

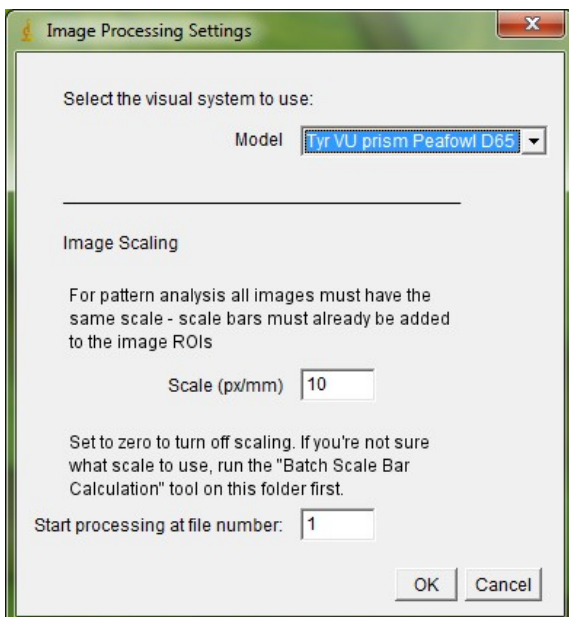


Image 23: Batch processing scaling and visual system options.

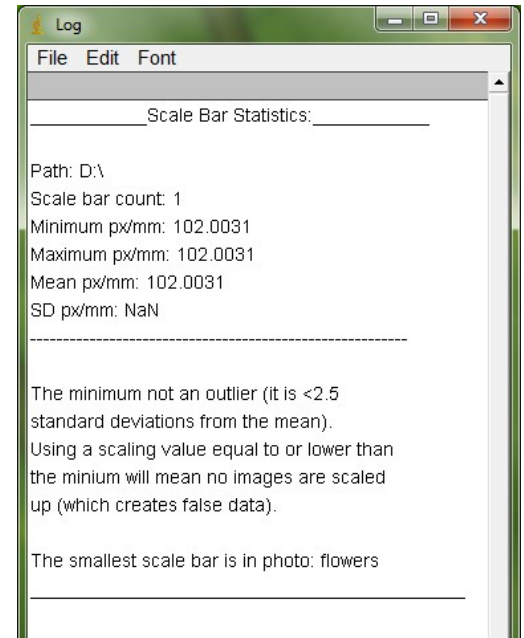
Once you have selected the directory a dialogue box will ask you about visual system and scaling options (Image 23).

**Model:** Select the model for converting from camera colours to your desired visual system. If your camera/visual system model is not available and you know the spectral sensitivities of your camera you can generate a mapping function (see page 31). If you do not know your camera's spectral sensitivities select 'none' to measure the

normalised camera values (the images will still be objective).

**Add human luminance channel:** Tick this box if you're working with human vision. This runs a script that adds human luminance as  $(LWS+MWS)/2$  the average of red and green, thought to code for the human luminance channel.

**Image Scaling:** If you are not doing pattern analysis and have not selected scale bars in your images set this to zero and no image scaling will be performed. For pattern analysis all your images need to have the same pixels per unit length. Every series of photos will have its own ideal scale that depends on the size of the samples you're working with. Generally you want to set this so that it reduces the size of all your images slightly (enlarging the images would create false data). The "Batch Scale Bar Calculation" tool will tell you which scale bar is the smallest, and whether this minimum value looks too small compared to your other images (just to make sure you are not using an anomalously small one). Try using this minimum value (or round down to an integer) and if the processing is going too slowly choose a smaller number (halve the number and it should go about four times faster). To use this tool go: **Plugins> Multispectral Imaging> Image Analysis> Batch Scale Bar Calculation**. All scaling is performed as bilinear interpolation to minimise the creation of artefacts.



*Image 24: Output from the Batch scale bar calculation tool. The minimum scale reported here is 102, so using this scaling value would be sensible, or round it down to 100 (do not round up).*

**Start processing after file number:** While processing very large datasets it can be annoying if it stops half way through (e.g. due to a power cut or crash). The batch processing records measurements as it proceeds, so there is no need to start processing from the start again, just put the image number to start at in this box to restart from where it got to. You can look in the output files in the same directory to see where it got to previously.

**Click OK** to proceed. The first multispectral image in the selected directory will be loaded so that the script knows what receptor channels are available.




Image 25: Pattern & luminance measurement dialogue.

Next, a dialogue will open asking what measurements should be made.

**Luminance channel:** The luminance channel is used for pattern measurements, but this will vary between different taxa. For human vision this would be “Lum”, added if you ticked the box in the previous dialogue. For birds this will be the double cones (“dbl”, as shown in Image 25). For dichromatic mammals just select the “lw” channel as this is most likely to encode luminance.

**Rescale Image:** This is a second scaling option available if you want to measure the pattern at a different scale to that selected in the previous dialogue, or want to scale all images uniformly (rather than based on a scale bar). In general this should always be set to 1 (off).

**Start size:** Specify the smallest scale to start bandpass filtering at. It is not possible to measure wavelengths lower than 2px in size, so this is a sensible start value. Set this to zero

to switch off pattern analysis.

**End size:** Pattern analysis will stop at this size. As a rule, this size should be equal to or smaller than your smallest sample's longest dimension. E.g. of your smallest sample is 50mm long and your px/mm is 20, this sample will be  $50 \times 20 = 1000$  pixels long. So the end size should be equal to or smaller than 1000px.

**Step Size:** Specify what scaling number to use for increasing the scale from the start size to the end size.

**Step multiplier:** Specify whether the scale should increase exponentially or linearly. The example in Image 25 will measure 13 levels between 2 and 128 pixels, increas-

ing as a multiple of 1.41; .i.e. 2, 2.8, 4, 5.7, 8, etc... Exponential steps generally make more sense than linear given the Gaussian filtering involved with the Fourier band-pass.

**Output pattern spectrum:** Tick this box if you want to measure pairwise pattern differences between samples. It saves the pattern data across all spatial scales, otherwise only the descriptive statistics are measured.

**Luminance bands:** This saves a basic luminance histogram, used for pairwise comparisons of multi-modal luminance distributions. For example, if you are measuring samples that have discrete patches of luminance, such as zebra stripes that are either black or white it does not always make sense to use the mean (grey in this case) to compare zebra to backgrounds as this grey level is not actually found on the zebra. You can set how many levels to save, around 20 to 100 might be sensible, though this number is somewhat arbitrary. Larger numbers of pixels could support a higher number of bins. If your histograms are always skewed to dark pixels, you could apply the log or square transform here.

**Combine ROIs with prefix:** Here you can specify any ROI groupings. E.g. entering "f,g" as in Image 25 will group together the petals of each respective flower from the example in Image 17, treating all objects labelled "F" as one object, and all

"G" objects as one. You can leave this box empty out to measure every ROI individually (e.g. this makes it easy to do inter-flower measures and intra-flower measures).



Image 26: Examples of pattern energy map outputs.

**Output energy maps:** Outputting the energy maps saves tiff images of the pattern analysis across all spatial scales (e.g. Image 26). This is mostly just for checking everything is working correctly – e.g. just run it on a sub-sample of files to check it is grouping objects together correctly. These are quite large files that you can delete once you have had a look at them – they are not required for any measurements.

**Click OK** and it will measure all .mspec images in the selected folder with the settings you have chosen.

## Batch image analysis results

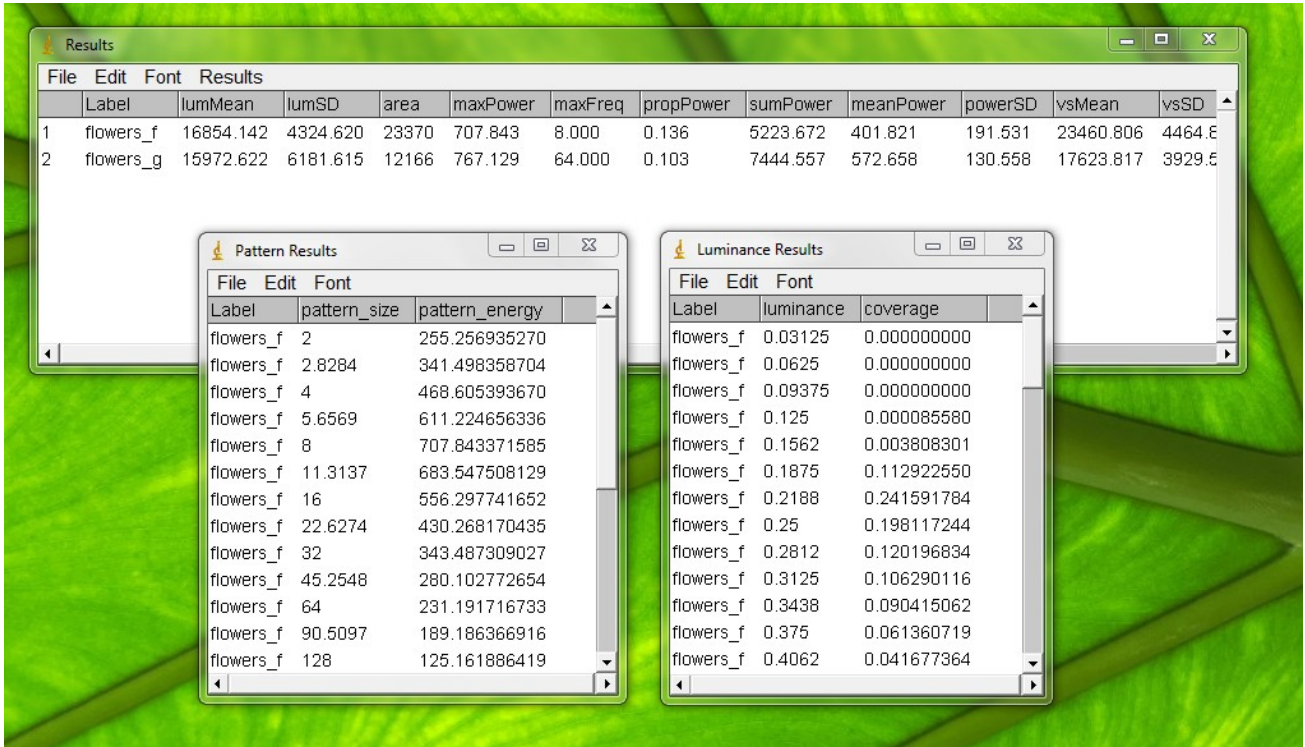


Image 27: Results from batch image analysis. These results are also saved in spreadsheet files in the image directory.

The results of batch image processing are shown in Image 27. The main colour measurements (mean and standard deviation for each ROI in each colour channel) and pattern descriptive data are saved to the results window. Note that luminance standard deviation is a measure of contrast. Depending on whether you opted to output pattern spectrum data and create luminance histograms these data will be shown in their own windows. All data are saved in spreadsheets in the image directory, and the data are saved as the batch processing progresses so the process can be restarted picking up from where it stopped.

## Granularity pattern analysis introduction

We provide tools for performing a pattern analysis based on Fast Fourier bandpass filtering, often called a granularity analysis. This form of analysis is increasingly widely used to measure animal markings (Godfrey et al., 1987; Stoddard and Stevens, 2010), and is loosely based on our understanding of low-level neuro-



physiological image processing in numerous vertebrates and invertebrates. Each image is filtered at multiple spatial frequency scales, and the “energy” at each scale is measured as the standard deviation of the filtered pixel values. The pattern energy of non-rectangular regions of interest are measured by extracting the selection so that all surrounding image information is removed (i.e. the selection area is measured against a black background). This image is duplicated and the selection area is filled with the mean measured pixel value (to remove all pattern information inside the selection area). Identical bandpass filtering is performed on both images and the difference between the images is calculated before measuring its energy for a shape-independent measure of pattern.

To perform pattern analysis using the batch measurement tool one must first select which channel to use for pattern analysis. Pattern processing in humans is thought to rely on the luminance channel (which is the combination of LW and MW sensitivities) and double cones in birds (Osorio and Vorobyev, 2005). If images are not being converted to cone catch quanta then the green channel is recommended (Spottiswoode and Stevens, 2010), or a combination of red and green (which is available as an automated option). Next, the desired measurement scales must be selected in pixels, along with the desired incrementation scale (e.g. linear increases of two pixels would provide measurements at 2, 4, 6, 8 pixels, and so on). A multiplier can be used instead, for example multiplying by two to yield measurements at 2, 4, 8, 16 pixels and beyond. A size no larger than the scaled image dimensions should be used to cover the entire available range. The number of scale increments should be judged based on processing speed, although increasing the number of scales measured will yield progressively less additional information. We find that increasing the scale from 2 by a multiple of  $\sqrt{2}$  up to the largest size available in the scaled images produces good results in many situations.

Summary statistics (Chiao et al., 2009; Stoddard and Stevens, 2010) of the pattern analysis are saved for each region of interest (or pooled regions). These include the maximum frequency (the spatial frequency with the highest energy; i.e. corresponding to the dominant marking size), the maximum energy (the energy at the maximum frequency), summed energy (the energy summed across all scales, or amplitude; a measure of pattern contrast), proportion energy (the maximum energy divided by the summed energy; a measure of pattern diversity, or how much one pattern size dominates), mean energy and energy standard deviation. In addition the energy spectra (the raw measurement values at each scale) can be output and used for pattern difference calculations (see below). Pattern maps can also be created – these are false colour images of the filtering at each scale –

produced for subjectively visualising the process.

Animal markings and natural scenes often have pattern energy spectra with more than one peak frequency, meaning the pattern descriptive statistics (above) can arbitrarily jump between peaks with similar energy levels in different samples (Image 28). An alternative approach we recommend when directly comparing two samples, rather than deriving intrinsic measurements to each sample, is to calculate the absolute difference between two spectra (A and B) across the spatial scales measured s:

$$Pattern_{diff} = \sum_{s=\max}^{s=\min} \left| (A_s - B_s) \right|$$

This is equivalent to the area between the two curves in Image 28. Any two patterns with very similar amounts of energy across the spatial scales measured will produce low pattern difference values irrespective of the shape of their pattern spectra. These differences will rise as the spectra differ at any frequency. As with colour, these differences can be calculated automatically between regions of interest, within or between images from the pattern energy spectra (above) using the "Pattern and Luminance Distribution Difference Calculator". Luminance

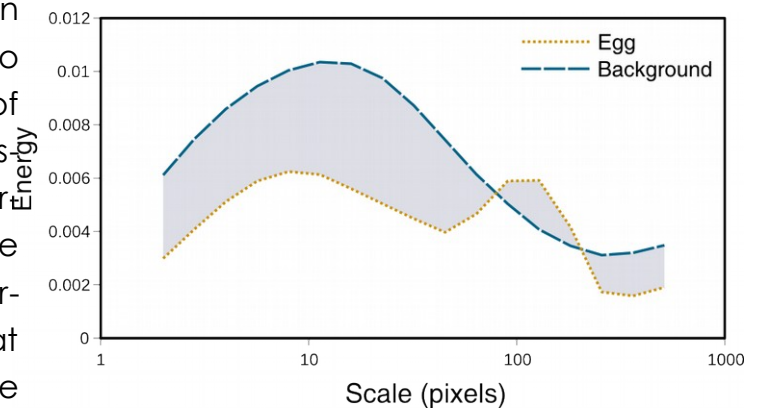


Image 28: Pattern spectra for an egg against its background. Note there are two peaks in spatial energy for the egg, implying it has both small spots and large spots. The peak frequency could therefore either be around 10 pixels or 100 pixels.

distribution differences are calculated similarly, summing the differences in the number of pixels in each luminance bin across the entire histogram. This method for comparing luminance is recommended when the regions or objects being measured do not have a normal distribution of luminance levels. For example, many animal patterns have discrete high and low luminance values (such as zebra stripes or egg maculation), in which case a mean value would not be appropriate. Comparing the luminance histograms overcomes these distribution problems.

## Calculating pairwise pattern and luminance differences

Descriptive pattern data such as the peak energy and peak spatial frequencies of pattern spectra can be useful for comparing treatments, but for pairwise comparisons when there is more than one peak (such as the egg in Image 28) it makes more sense to measure pattern difference as the area between the curves (in this case the grey area between the egg and its background). The same principle can be used for comparing multi-modal luminance histogram data.




Image 29: Pattern and luminance distribution difference calculator options.

The toolbox can calculate this for you. Open your pattern spectrum or luminance histogram data (e.g. from the .csv file saved by batch analysis) so that they are in the results window (close any results windows open already and go **File>open**, or just drag the file into ImageJ), then run **Plugins>Multispectral Imaging>Data Analysis>Pattern and Luminance Distribution Difference Calculator**. The dialogue box will ask for a few options Image 29:

**Data:** Specify the column name that contains the pattern energy or luminance values.

**Compare this region:** Select one of the regions IDs from the ROI coding specified in the image.

**To this region:** The region ID to compare to. This can be the same, in which case it will only be able to compare between photographs.

**Region comparison.** “Within photo” would compare the first region ID to the second region ID, comparing flower “f” to flower “g” in this example. Select “Between photos” to compare flower “f” to any instances of flower “g” across all

photos (this will compare within photo if applicable as well).

**Click OK**, it will generate the comparisons and create a table of results. You can then save this table as a .csv file, or copy and paste the data into a spreadsheet.

## Colour differences introduction

There are various methods available for comparing colours, such as hue calculations, principal components analysis, or opponent channel responses. A common method for comparing colours in non-human visual systems, in terms of whether two objects are likely to be discriminable or not, is Vorobyev and Osorio's (1998) receptor noise model which calculates "just noticeable difference" (JND) values. This predicts whether two similar colours are likely to be discriminable based on the signal to noise ratios of each channel to di-, tri- and tetrachromatic visual systems. JND values less than one are indiscriminable, and values above ~3 are discriminable under good lighting conditions (Siddiqi et al., 2004). The toolbox implementation of the model only applies to photopic vision (daylight lighting conditions), and uses the standard log version of the model. Luminance JNDs follow Siddiqi et al. (2004). Some authors argue that JNDs are not a suitable metric to use for comparing highly dissimilar colours (e.g. see Endler & Mielke, 2005), although the performance of different colour space models is an area that requires further investigation. JNDs are currently therefore intended for comparisons between similar colours.

## Calculating pairwise colour and luminance discrimination values

"Just noticeable differences" (JNDs) tell you whether a given animal is likely to be able to distinguish two colours (or levels of luminance) under ideal viewing conditions (see above). To calculate colour or luminance JND differences open the results .CSV file containing your cone-catch results (so it's in the "Results" window of ImageJ) and go: **Plugins> Multispectral Imaging>Data Analysis>Calculate colour JNDs** (see Image 30). This calculation uses the common log-transformed model, and should only be used for photopic (daytime) vision with plenty of light (Siddiqi et al., 2004; Vorobyev and Osorio, 1998). This tool only compares the mean colour values for each region you have measured. If there are lots of very different colours in those regions this method is not suitable. If the sample you are measuring has lots of discrete colours then these regions should be measured in different ROIs,

or an image quantisation method should be used.

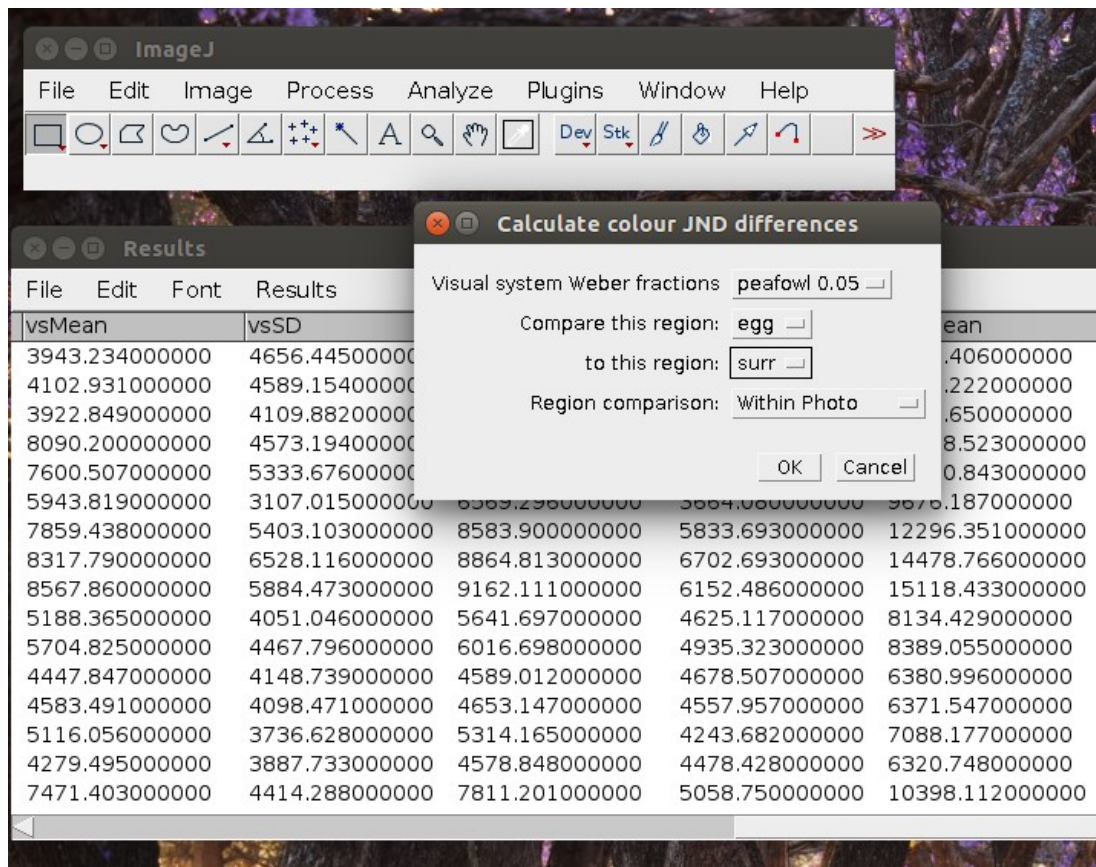


Image 30: Calculate pairwise colour differences using JNDs.

**Visual system Weber fractions:** Select your visual system's Weber fractions. These are often just calculated from relative cone abundance values – you can add your own Weber fractions by making a new .txt file in plugins/Multispectral Imaging/weberFractions. Examples are included in the toolbox for common model visual systems with global Weber fractions of 0.05 and 0.02.

**Compare this region:** First region to use in the comparison.

**To this region:** Second region to use in the comparison (can be the same as the first region for between-photo comparisons).

**Region comparison:** Specify whether the comparison should be within or between regions and photos.

**Click OK**, and the comparisons will be calculated. Save the data as a .csv file, or copy and paste into a spreadsheet.

---

# PRESENTING DATA

---

## Statistics

Once you have generated colour, pattern, or luminance metrics the data can be analysed in any suitable statistics package. When using analyses that assume normality of error structure such as ANOVAs, GLMs, GLMMs etc these variables will often benefit from a log transformation. Luminance distribution difference values are bounded at 2, i.e. the difference between the normalised histograms have a maximum of 2 if there's no overlap between the sample's luminance values at all. Where these numbers appear to be bounded, a solution is to divide all the values by 2 (so the maximum is 1), and then apply a logit transform:

$$\text{logit}(p) = \log\left(\frac{p}{1-p}\right)$$

This should make the luminance difference data normally distributed even if they are near the upper and lower bounds.

## Reporting your methods

When publishing findings based on image analysis it is important to report the equipment and methods used in sufficient detail. Here is a list of the things that you should include so that your study could be replicated:

- Camera model, reporting whether it was converted to full spectrum.
- Lens model (if it was a zoom lens, what focal length was used?).
- Filters, models and approximate transmission characteristics (e.g. "Baader Venus-U filter transmitting from ~320 to 380nm").
- Light source or lighting regime (e.g. "direct sunlight, not within 1 hour of sunrise or sunset", or "Iwasaki eyeColor arc lamp with its UV filter removed").
- Grey standard reflectance values and model or material. If a direct/point light source was used, what angle was the grey standard relative to the illuminant (e.g. "the grey standard was always angled towards the illuminant" or "the grey standard was level with the plane of the sample" (see page 10)?

- How were regions of interest selected? If there is some ambiguity in a region's outline how was this accounted for? In this case consider having the regions selected by someone who is blind to the hypothesis and treatment being selected.
- If image scaling was used what was the scaling value, and what was the average size of the sample being measured (e.g. "images were scaled to 20 px/mm, the average sample was  $52 \pm 3.5$  mm in length").
- If you converted your images to cone-catch values report where the camera spectral sensitivity functions were sourced, or how they were measured, what database of natural spectra were used for the modelling, and where the animal's receptor spectral sensitivities came from. The  $R^2$  values for the model fit to natural spectra could also be reported (these numbers can be retrieved by opening the .java file for your model in imagej/plugins/Cone Models/).
- If you are measuring JNDs report the Weber fraction used for the most abundant receptor channel (this will normally be 0.05 or 0.02), and cone proportions.
- If you are measuring pattern (descriptive statistics or pairwise pattern differences) what spatial frequencies were measured?
- If you compared luminance distribution differences how many bins were used?
- Will your image data be made available in a public archive? RAW files should be made available rather than any other image format to ensure compatibility with these tools. The use of RAW files also ensures no non-linear processing or destructive processing has been performed, and maintains the meta-data of each image. We would also recommend uploading the associated .mspec files and ROI files for each image as these are very small files that enable quick data analysis.

See White et al. (2015) for further information on repeatability in colour measurements.

## Creating colour or false-colour images for presentation

Linear images are essential when measuring reflectance values, but they normally look dark and dull on your monitor. This is because the dynamic range of cameras is far higher than that of your viewing media (e.g. monitor or printed on paper), meaning your camera can detect a far higher difference between the lightest

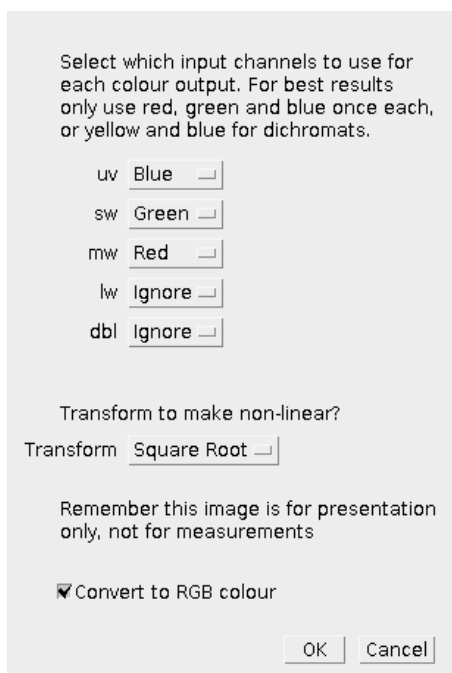


and darkest parts of a scene than your monitor or printer can recreate. To account for this photos are non-linearly transformed, normally with a Gamma curve to try to squeeze the image's dynamic range into a smaller range.

When presenting examples of your photographs it is acceptable to transform the images so that they are non-linear for display on monitors and in print, though it is good practice to make it clear what transform has been applied.


Presenting a tetrachromatic image is obviously impossible for human viewing, so instead one can present a false colour or pseudo-colour image where a subset of the channels are arranged into an RGB stack (Image 32).

1. Open your .mspec image: **Plugins>Multispectral Imaging>Load Multispectral Image**.
2. Convert your image to the cone-catch values of your species (ignore this step if you just want to show objective camera values): **Plugins>Multispectral Imaging>Image Analysis>Convert to Cone-Catch**, and select your camera/visual system model when prompted.
3. At this stage the image will be a grey-scale 32-bit multi-channel stack of cone-catch values or camera pixel values. Run **Plugins>Multispectral Imaging>Tools>Make Presentation Image**.



*Image 31: Select channels to convert to a colour image suitable for human viewing. This example will create a false-colour UV image from a bluetit cone-catch image.*

4. Select which cone-catch channels or image channels you want to convert into which colour for viewing (see Image 31 for a bluetit example). Ideally select red, green and blue once each. Or, if you have a dichromatic animal select yellow and blue only. Leave all other unwanted channels as "Ignore".
5. **Transform:** Applying a square-root transform will reduce the image contrast and make the darker regions more visible. Optionally leave this as linear by selecting "none".
6. **Convert to RGB Colour:** This will output a "normal" RGB image that you can save as a JPG or TIF, and use in other software packages or for printing.
7. When presenting this image report which channels are being displayed and that the linear values were square-root transformed to optimise it for displaying. See Image 32 for an example.



*Image 32: Example of Gallotia galloti male and female lizards. On the left is a "normal" human CIE XYZ image, on the right is a false-colour bluetit image showing MW, SW and UV channels. Both images are square-root transformed from linear cone-catch images for display purposes. The lizard's blue cheek on the right-hand image demonstrates that this patch has very high UV reflectance relative to MW and SW.*

---

# REFERENCES

---

- Chakrabarti, A., Scharstein, D., and Zickler, T. (2009). An Empirical Camera Model for Internet Color Vision. In *BMVC*, (Citeseer), p. 4.
- Chiao, C.-C., Chubb, C., Buresch, K., Siemann, L., and Hanlon, R.T. (2009). The scaling effects of substrate texture on camouflage patterning in cuttlefish. *Vision Res.* 49, 1647–1656.
- Endler, J.A., and Mielke, P.W. (2005). Comparing entire colour patterns as birds see them. *Biol. J. Linn. Soc.* 86, 405–431.
- Godfrey, D., Lythgoe, J.N., and Rumball, D.A. (1987). Zebra stripes and tiger stripes: the spatial frequency distribution of the pattern compared to that of the background is significant in display and crypsis. *Biol. J. Linn. Soc.* 32, 427–433.
- Hong, G., Luo, M.R., and Rhodes, P.A. (2001). A study of digital camera colorimetric characterisation based on polynomial modeling. *Color Res. Appl.* 26, 76–84.
- Lovell, P.G., Tolhurst, D.J., Párraga, C.A., Baddeley, R., Leonards, U., Troscianko, J., and Troscianko, T. (2005). Stability of the color-opponent signals under changes of illuminant in natural scenes. *JOSA A* 22, 2060–2071.
- Osorio, D., and Vorobyev, M. (2005). Photoreceptor spectral sensitivities in terrestrial animals: adaptations for luminance and colour vision. *Proc. R. Soc. B Biol. Sci.* 272, 1745–1752.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing* (Vienna, Austria: R Foundation for Statistical Computing).
- Schneider, C.A., Rasband, W.S., and Eliceiri, K.W. (2012). NIH Image to ImageJ: 25 years of image analysis. *Nat Methods* 9, 671–675.
- Siddiqi, A., Cronin, T.W., Loew, E.R., Vorobyev, M., and Summers, K. (2004). Interspecific and intraspecific views of color signals in the strawberry poison frog *Dendrobates pumilio*. *J. Exp. Biol.* 207, 2471–2485.
- Spottiswoode, C.N., and Stevens, M. (2010). Visual modeling shows that avian host parents use multiple visual cues in rejecting parasitic eggs. *Proc. Natl. Acad. Sci.* 107, 8672–8676.
- Stevens, M., Parraga, C.A., Cuthill, I.C., Partridge, J.C., and Troscianko, T.S. (2007). Using digital photography to study animal coloration. *Biol. J. Linn. Soc.* 90, 211–237.
- Stoddard, M.C., and Stevens, M. (2010). Pattern mimicry of host eggs by the common cuckoo, as seen through a bird's eye. *Proc. R. Soc. B Biol. Sci.* 277, 1387–1393.
- Troscianko, J. (2014). A simple tool for calculating egg shape, volume and surface area from digital images. *Ibis* 156, 874–878.
- Vorobyev, M., and Osorio, D. (1998). Receptor noise as a determinant of colour thresholds. *Proc. R. Soc. Lond. B Biol. Sci.* 265, 351–358.

Westland, S., Ripamonti, C., and Cheung, V. (2004). Characterisation of Cameras. *Comput. Colour Sci. Using MATLAB®* Second Ed. 143–157.

White, T.E., Dalrymple, R.L., Noble, D.W.A., O’Hanlon, J.C., Zurek, D.B., and Umbers, K.D.L. (2015). Reproducible research in the study of biological coloration. *Anim. Behav.* *106*, 51–57.