

Illumination Correction tutorial

I. Introduction

The **Correct Illumination - Calculate** and **Correct Illumination - Apply** modules are intended to compensate for the non-uniformities in illumination often present in microscopy images. It is not uncommon for the intensity within a fluorescence image to vary by more than two-fold across the field of view due to the optics of the microscope, imperfection in the slide or sensor bias. Correcting the illumination variation can be important both for accurate segmentation and for intensity measurements.

II. Creating an illumination function

The Correct Illumination - Calculate module is used to create the illumination function, an image that is representative of the overall pattern of illumination in the image or image set. Generally, it is necessary to create a different function for each set of imaging or sample preparation conditions because either can yield changes in the pattern of illumination. For example, the illumination pattern will change if you use a different staining reagent for a batch of images or change any components or settings in the optical path of the microscope. In fact, sometimes the illumination pattern can change throughout the course of the day as the microscope lamp changes temperature, or between different batches of what should be identical staining reagents. In our experience, a separate illumination function should be calculated for each separate plate in a multi-well plate experiment. Additionally, while the pattern of illumination may look very similar for each imaged wavelength, the absolute intensities will most likely vary between wavelengths (channels). Therefore, a separate illumination function should be prepared for each channel.

Deciding on a workflow: There are two workflow options available in CellProfiler to create and apply illumination functions for a project: (1) create and save illumination functions using one CellProfiler pipeline, then later retrieve and apply the illumination functions using a separate CellProfiler analysis pipeline; or (2) create and apply illumination functions within a single CellProfiler analysis pipeline. For experiments where a different illumination function will be calculated for each individual image, (2) is the better choice. For typical high-throughput experiments involving more than 100 images that will share an illumination function, (1) is more convenient because it allows you to readily inspect the illumination functions that are produced for quality control purposes, prior to applying them to analyze images. It offers the opportunity to tweak the pipeline that creates the illumination functions prior to analyzing images. As well, if creating the illumination function is time-consuming (i.e., many images are being processed), it is worthwhile to create the function in a pipeline separate from the one used for analysis, to make the analysis pipeline more efficient. Example pipelines following both workflows are available at www.cellprofiler.org/examples.shtml in the IlluminationCorrection example pipeline.

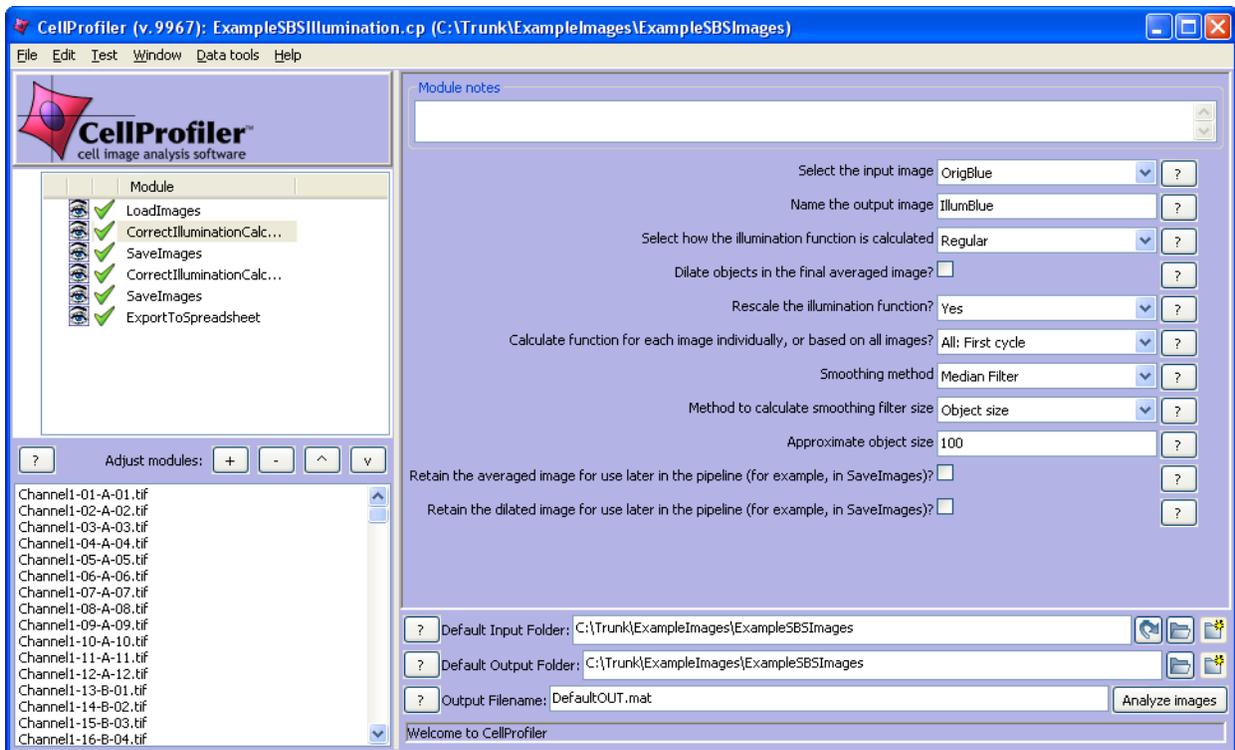


Figure 1: Screenshot of the CellProfiler interface for the Correct Illumination - Calculate module.

Configuring the Correct Illumination – Calculate module to create illumination functions: An example of the Correct Illumination - Calculate module is shown in Fig. 1. Some module settings only become visible based on your response to other settings, so only some of the available settings are shown here. For overall guidance on using the module, see the full help for the module by selecting it in the module panel at the top left of the CellProfiler window, then click the Click the Help button  under the module panel on the left. Click the Help buttons to the right of each module setting for a detailed description of that setting. If saving an illumination function for later retrieval, be sure to select '.mat' as the file format for saving since this format uses floating point pixel values, which is necessary for the mathematical operations involved in illumination correction.

III. Applying the illumination function

In the Correct Illumination - Apply module (some settings of which are shown in Fig. 2), you have the option to either divide or subtract the illumination function from the input image (the image to be corrected). If the illumination function was rescaled to be 1 or greater in Correct Illumination - Calculate, select *Divide*. Otherwise, if the function was not rescaled, select *Subtract*. See the help for the *Background vs. Regular* setting in the Correct Illumination - Calculate module for more guidance. In both cases, this should leave the image in the range of 0 to 1.

If the gradient across the final corrected image remains large, you may want to rescale the intensities after the Apply step by using the **Rescale Intensity** module on the final image. Since doing so will stretch the intensities in the final image, this function should be used carefully. Rescaling is not recommended when you intend to measure and compare intensities among rescaled images in the image set. It is also not recommended if the image set may contain images that have no objects; in such cases the rescaled image will set the brightest part of the blank background of the image to be very bright, thus confusing later modules such as Identify Primary Objects. See the help for the Rescale Intensity module for more information on the rescaling options available.

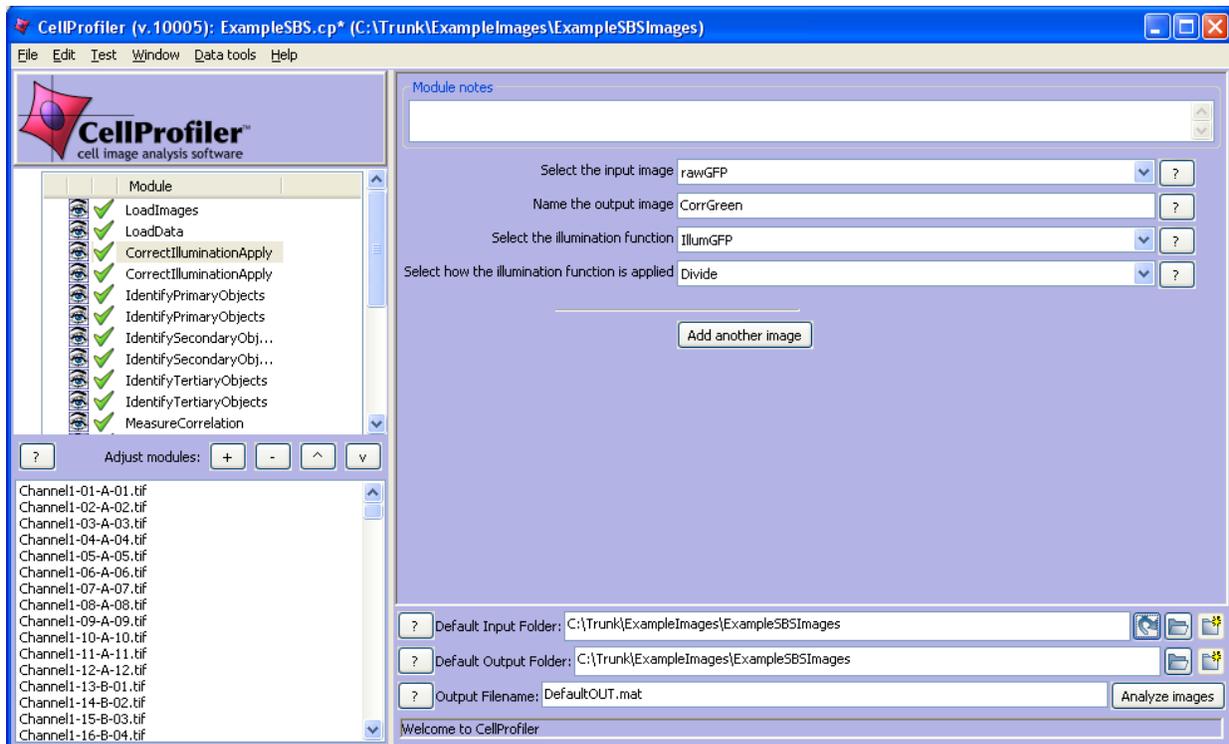


Figure 2: Screenshot of the CellProfiler interface for the Correct Illumination - Apply module.

IV. Inspecting the illumination function and its results

The quality of the Illumination correction depends entirely on the settings chosen for Correct Illumination - Calculate. Your choice as to what combinations of settings are appropriate will depend on the spatial arrangement of the features of interest in your cellular images and whether each image is likely to have a different illumination pattern or whether multiple images are likely to share the same pattern. In this section, we will inspect example illumination functions and assess whether they are appropriate.

Before proceeding: Please gain a basic understanding of the options within *Correct Illumination – Calculate and Correct Illumination – Apply* by reading the help for the module: select it in the module panel at the top left of the CellProfiler window, then click the *Click the Help* button under the module panel on the left.

Example 1: In this example from a collection of images of U2OS cells (Fig. 3):

1. The cells are uniformly distributed across the image (that is, they are not preferentially located in certain parts of the image).
2. They occupy most of the foreground.
3. There seems to be a complex illumination pattern where the illumination seems dimmer in the top left and possibly the bottom left and top right corners.
4. The image is part of a batch of images prepared under the same sample preparation and imaging conditions, all of which show a similar illumination pattern.

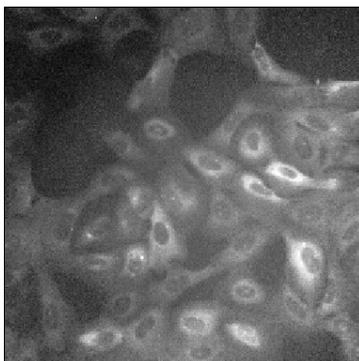


Figure 3: Raw grayscale image of U2OS cells.

This image set should be corrected using the *Regular* method, which tends to produce more accurate correction when there bright objects across much of the image. Because of (2), there is insufficient background in the image for the alternative *Background* method. Although other options might work, we will choose *Median* as the smoothing method because of (3) – the pattern appears to be rather complex as is often the case for fluorescence images. Because of (4), we should choose *All* (to calculate a function based on many images together) rather than *Each* (which calculates a different function for each image). If we were to incorrectly choose the *Regular* method upon *Each* image individually, we would obtain a poor illumination function that resembles the original distribution of the cells, showing dark regions precisely where cells are absent and bright regions where many cells are present (Fig. 3A). Since this function is more reflective of the extrinsic features of interest (the cells themselves), and not the intrinsic qualities of the acquisition system, this is not a desirable illumination function. However, because this image is part of a set produced by an automated microscope, it would be better to take advantage of the whole image set to produce a more robust illumination function that is less sensitive to variations in each particular image. It is thus preferable to use *All* instead of *Each*, averaging together many images prior to smoothing. The result is an ensemble illumination function that is more representative of the fluorescence variation intrinsic to all images captured during the experiment (Fig. 4B). Even so, the illumination pattern still has some foreground

variations originating from the cell distribution instead of the imaging system. Increasing the smoothing filter size yields an illumination function that (Fig. 4C).

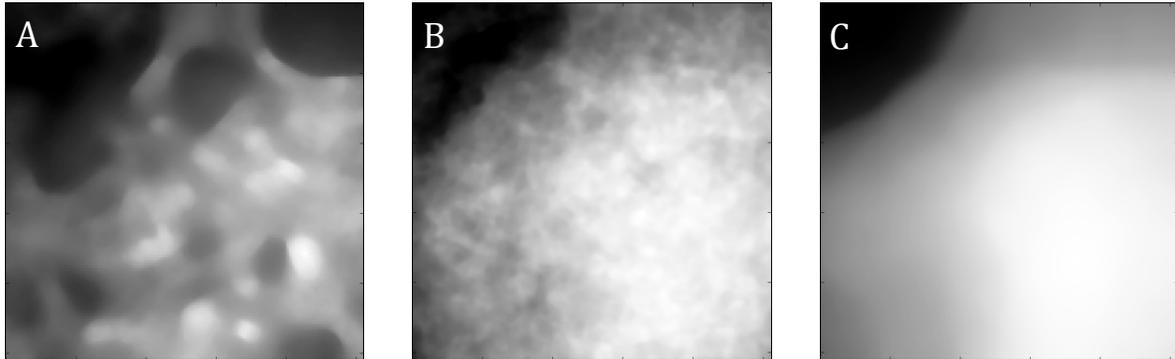


Figure 4: Illumination functions produced from the image set containing Fig. 3 (A) Using *Regular + Each + Median* filtering on the raw image in Fig. 3. (B) Using *Regular + All + Median* filtering on the full image set. (C) Using *Regular + All + Median* filtering on the full image set with a larger smoothing filter size.

Figure 5 shows the results of the illumination correction to this example. After applying the illumination function in Fig. 4C to the original image using *Division* in *Correct Illumination - Apply*, we get the following result in Fig. 5B. The cells are much more even in intensity, with less variation between the outer edges and the center of the image. Caution should be used to make sure the images are not over-corrected – that is, to make sure that real variation in the brightness of cells is not removed using illumination correction techniques, disturbing the measurements made from the cells.

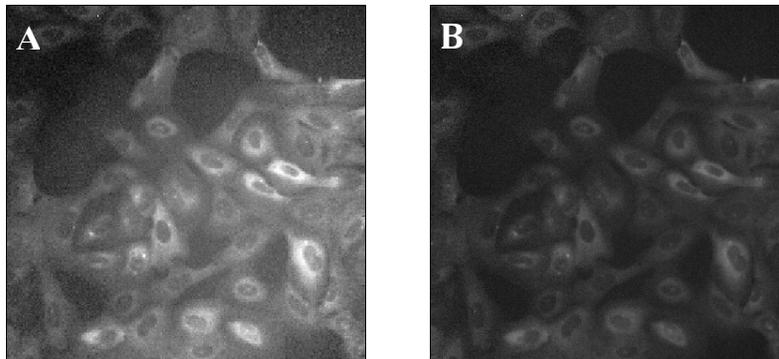


Figure 5: The original image of cells (A) corrected for illumination variation (B).

Example 2: In this example of an image of nuclei stained with histone-2B cherry and containing a large image artifact (Fig. 6A):

1. The organisms are sparsely distributed.
2. The background of the image seems to show the pattern of illumination.
3. The illumination pattern varies substantially between different images in the set (not shown).

An attempt to identify the nuclei in the original image will result in detection of the artifact in addition to the nuclei present (Fig. 6B). Because of (1) and (2), the *Background* method is the more appropriate option for correcting the image as opposed to *Regular*. Because of

(3), we should select *Each* rather than *All*, since each image will need its own illumination correction function.

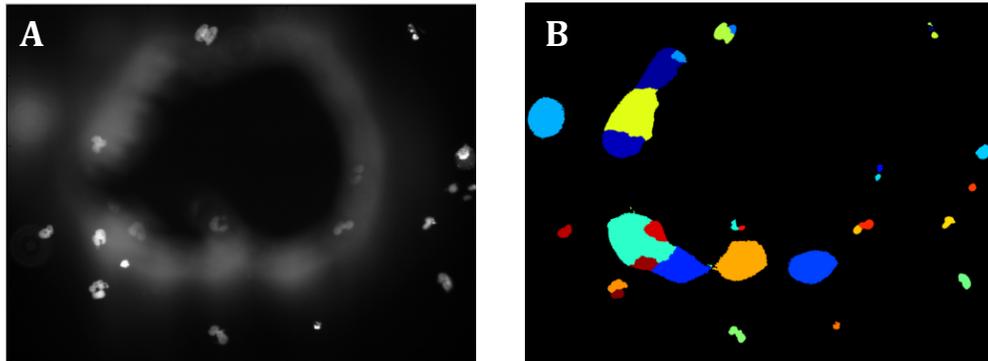


Figure 6: (A) Original image of nuclei corrupted with an intensity artifact. (B). Downstream effect on nuclei identification.

Using *Background + Each + Median* as in Example 1 yields the illumination functions shown in Fig. 7A and B. The illumination function in Fig. 7A, however, used a block size that was too small in Correct Illumination - Calculate, so while the background intensity distribution is visible, some foreground pixels were included in the function. Upon inspection, it is easy to see the nuclei in the illumination function, which is undesirable – portions of the nuclei will be improperly removed from the image if this illumination function is applied. Using *Background + Each + Median* with a larger block size yields an illumination function which better reflects the actual background intensity distribution (Fig. 7B).

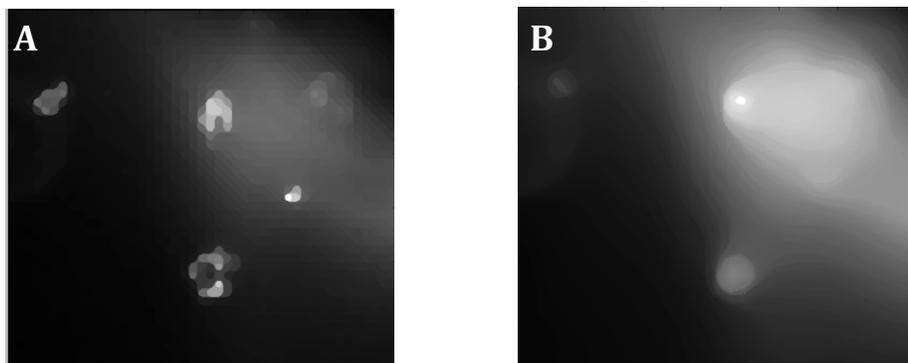


Figure 7: A zoomed-in view of Illumination functions produced from the image in Figure 6. (A) Using *Regular + Each + Median* filtering. (B) Same as (A) with a larger block size.

Figure 8A and B depict the corrected image with the results of nuclei identification. With the artifact removed, the nuclei are now segmented much more accurately.

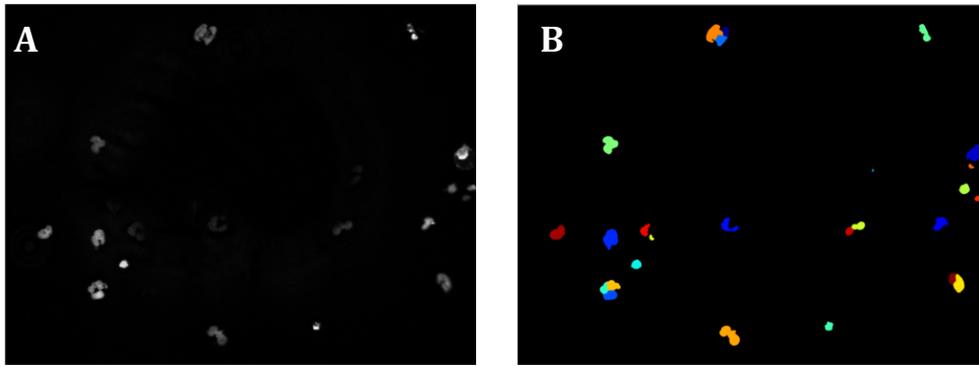


Figure 8: (A) Corrected image of nuclei. (B). Downstream effect on nuclei identification.

Example 3: In this example of *C. elegans* nematodes contained in a well (Fig. 9):

1. The organisms are sparsely distributed.
2. The background of the image (white portion in Fig. 9A) seems to show the pattern of illumination.
3. The illumination pattern varies between different images in the set (not shown).
4. The pattern tends to be a very simple one that is bright towards the middle and dim towards the edges of the well. This pattern is typical for brightfield images.

Because of (1) and (2), the *Background* method is the more appropriate option rather than *Regular*. Because of (3), we should select *Each* rather than *All*, since each image will need its own illumination correction function. In this instance, two pre-processing steps are performed prior to illumination correction. First, the image is inverted in intensity using the **Image Math** module (Fig. 9B), since the *Background* method assumes that the background is comprised of low intensity pixels. Second, the region outside the well is masked out using the **Mask Image** module in order to restrict the area of interest to the well interior. Without this method, Correct Illumination – Calculate will assume that the illumination was extremely dim at the edges of the image and attempt to correct the dark regions, resulting in a very skewed illumination pattern near the edges of the well. We choose *Background + Each + Fit Polynomial* due to (4) from the list above – the pattern appears to be simple as is often the case for brightfield images. This yields the illumination function shown in Fig. 9C.

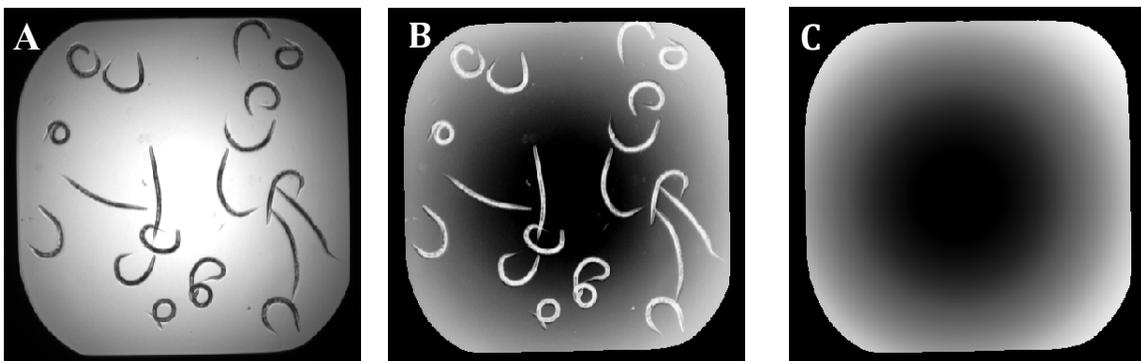


Figure 9: (A) Raw image of nematodes in a well in grayscale. (B) The image in (A) inverted in intensity. (C) Illumination functions produced from (B) using *Background + Each + Fit Polynomial*.

However, another correction method is available which yields similar results as well as being easier to use. The *Regular + Each + Convex Hull* method is optimized for transmitted light images with illumination patterns similar to the one illustrated in this image. In brief, the method effectively “erases” dark objects from the light background, which has the following advantages over the *Background* method above:

1. There are no requirements regarding the distribution of organisms in the image.
2. No inversion of pixel intensities (as in Fig. 9B) is needed since the method assumes the image consists of a dark foreground on a light background.

The image correction is applied by dividing the original image by the resultant illumination function i.e., using *Divide* in Correct Illumination – Apply.

This method produces the illumination correction function shown in Fig. 10B. Applying this illumination function to the original image yields a corrected image in which the background illumination within the well is absent (Fig. 10C).

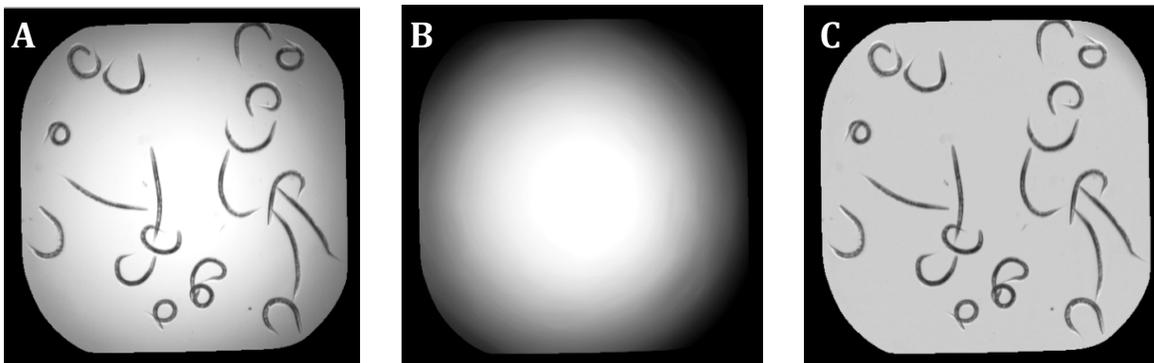


Figure 10: Images of the raw image (A), the illumination correction function (B) and the corrected image (C).