Computer Vision (600.461/600.661)
Project 1: Colorizing the Prokudin-Gorskii photo collection of images of the Russian Empire

Instructor: René Vidal

Due 11/04/2014, 11.59PM Eastern

Background. Sergei Mikhailovich Prokudin-Gorskii (1863-1944) was a man well ahead of his time. Convinced, as early as 1907, that color photography was the wave of the future, he won the Tsar's special permission to travel across the vast Russian Empire and take color photographs of everything he saw. And he really photographed *everything*: people, buildings, landscapes, railroads, bridges, ..., thousands of color pictures! His idea was simple: record three exposures of every scene onto a glass plate using a red, a green, and a blue filter. Never mind that there was no way to print color photographs until much later – he envisioned special projectors to be installed in "multimedia" classrooms all across Russia where the children would be able to learn about their vast country. Alas, his plans never materialized: he left Russia in 1918, right after the revolution, never to return again. Luckily, his RGB glass plate negatives, capturing the last years of the Russian Empire, survived and were purchased in 1948 by the Library of Congress (LoC). The LoC has recently digitized the negatives and made them available on-line at http://www.loc.gov/exhibits/empire/.

Overview. The goal of this assignment is to take the digitized Prokudin-Gorskii glass plate images such the one shown in Figure 1, and, using techniques learned in class, automatically produce a color image with as few visual

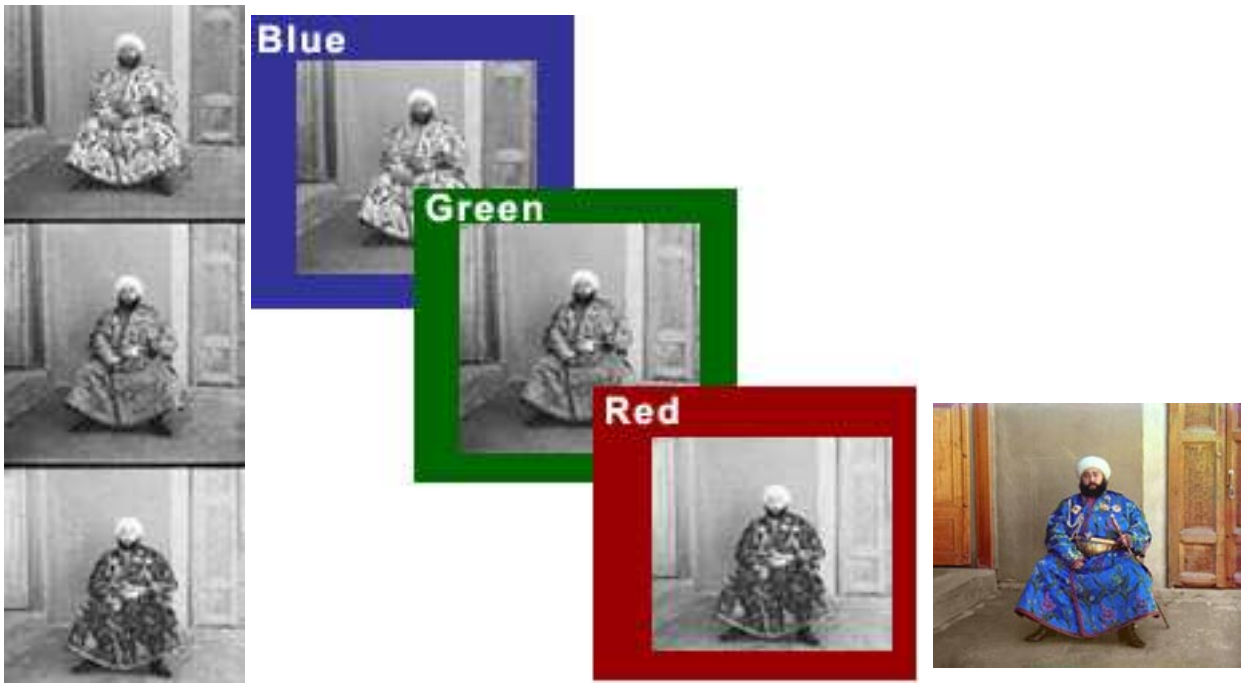

Figure 1: An example of the images taken by Prokudin-Gorskii, and the colored image you are expected to produce.

artifacts as possible. In order to do this, you will need to extract the three color channel images, place them on top of each other, and finally align them so that they form a single RGB color image. Recall that the methods discussed in class were designed to align two images. Here, you should describe how such methods can be extended to align three images. A few of the digitized glass plate images (both hi-res and low-res versions) can be found here. Note that the filter order from top to bottom is BGR, not RGB!. Note also that the images to be matched do not actually have the same brightness values (they are different color channels), so a clever metric for matching might be needed.

We will assume that a simple 2D translational or 2D affine model is sufficient for proper alignment. However, the full-size glass plate images are very large, so your alignment procedure will need to be relatively fast and efficient (e.g., by using an image pyramid). Finally, you should compare your results with those of the VLFeat software available at http://www.vlfeat.org. For this purpose, go through the tutorial http://www.vlfeat.org/overview/sift.html and learn how to extract, visualize and match SIFT features. Then modify the code in http://www.vlfeat.org/applications/sift-mosaic-code.html to align the three color channel images and produce a single color image.

Bells & Whistles. Although the color images resulting from this automatic procedure will often look strikingly real, they are still a far cry from the manually restored versions available on the LoC website and from other professional photographers (for example, check out this wonderful site: http://www.gridenko.com/pg/). Of course, each such photograph takes days of painstaking Photoshop work, adjusting the color levels, removing the blemishes, adding contrast, etc. As extra credit, you can try to make some of these adjustments automatically, without the human in the loop. Feel free to come up with your own approaches. There is no right answer here: just try out things and see what works. To see how well you are doing, you can compare your results with those at http://www.cs.cmu.edu/~dellaert/aligned/.

What to submit? Apart from your well documented code, you need to write a few paragraphs describing how you aligned your images, what metric you used and what variations you tried. Show your results on all images that were provided, plus a few others of your own choosing from the LoC collection.

[Exercise taken from the course "Rendering and Image Processing" offered at CMU by Alexai Efros.]