Project 2: SIFT Local Feature Matching

CS 6476

Spring 2021

Brief

- Due: Feb 19, 2021 11:59PM
- Project materials including report template: proj2.zip
- Additional scenes to test on extra_data.zip
- Hand-in: through Gradescope
- Required files: `<your_gt_username>.zip`, `<your_gt_username>_proj2.pdf`

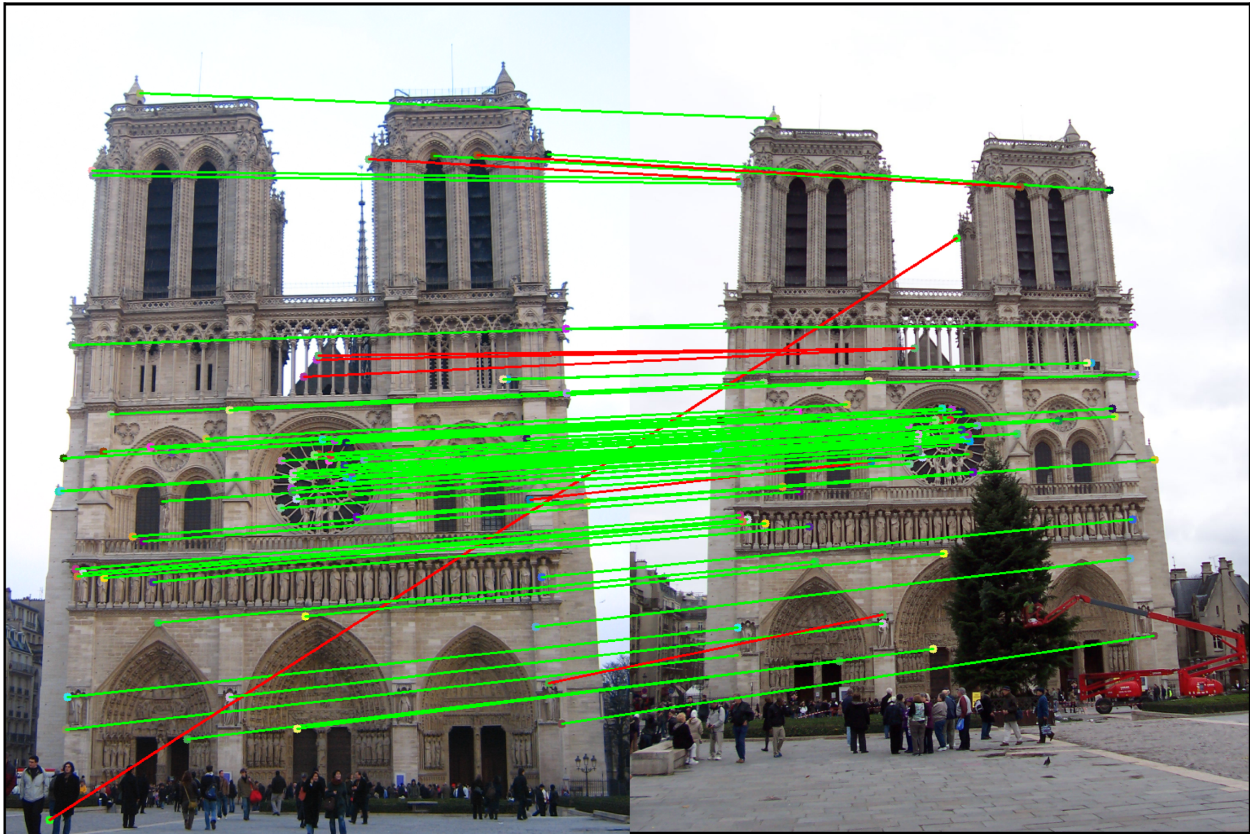

Figure 1: The top 100 most confident local feature matches from a baseline implementation of project 2. In this case, 89 were correct (lines shown in green), and 11 were incorrect (lines shown in red).

Overview

The goal of this assignment is to create a local feature matching algorithm using techniques described in Szeliski chapter 7.1. The pipeline we suggest is a simplified version of the famous SIFT pipeline. The matching pipeline is intended to work for *instance-level* matching – multiple views of the same physical scene.

Setup

1. Install Miniconda. It doesn't matter whether you use Python 2 or 3 because we will create our own environment that uses python3 anyways.
2. Download and extract the project starter code.
3. Create a conda environment using the appropriate command. On Windows, open the installed "Conda prompt" to run the command. On MacOS and Linux, you can just use a terminal window to run the command, Modify the command based on your OS (`linux`, `mac`, or `win`): `conda env create -f proj2_env_<OS>.yml`
4. This will create an environment named "cs6476_proj2". Activate it using the Windows command, `activate cs6476_proj2` or the MacOS / Linux command, `conda activate cs6476_proj2` or `source activate cs6476_proj2`
5. Install the project package, by running `pip install -e .` inside the repo folder. This might be unnecessary for every project, but is good practice when setting up a new `conda` environment that may have `pip` requirements.
6. Run the notebook using `jupyter notebook ./proj2_code/proj2.ipynb`
7. After implementing all functions, ensure that all sanity checks are passing by running `pytest proj2_unit_tests` inside the repo folder.
8. Generate the zip folder for the code portion of your submission once you've finished the project using `python zip_submission.py --gt_username <your_gt_username>`

Details

For this project, you need to implement the three major steps of a local feature matching algorithm (detecting interest points, creating local feature descriptors, and matching feature vectors). We'll implement two versions of the local feature descriptor, and the code is organized as follows:

- Interest point detection in `part1_harris_corner.py` (see Szeliski 7.1.1)
- Local feature description with a simple normalized patch feature in `part2_patch_descriptor.py` (see Szeliski 7.1.2)
- Feature matching in `part3_feature_matching.py` (see Szeliski 7.1.3)
- Local feature description with the SIFT feature in `part4_sift_descriptor.py` (see Szeliski 7.1.2)

1 Interest point detection (`part1_harris_corner.py`)

You will implement the Harris corner detection as described in the lecture materials and Szeliski 7.1.1.

The auto-correlation matrix A can be computed as (Equation 7.8 of book, p. 404)

$$A = w * \begin{bmatrix} I_x^2 & I_xI_y \\ I_xI_y & I_y^2 \end{bmatrix} = w * \begin{bmatrix} I_x \\ I_y \end{bmatrix} \begin{bmatrix} I_x & I_y \end{bmatrix} \tag{1}$$

where we have replaced the weighted summations with discrete convolutions with the weighting kernel w (Equation 7.9, p. 405).

The Harris corner score R is derived from the auto-correlation matrix A as:

$$R = \det(A) - \alpha \cdot \operatorname{trace}(A)^2 \tag{2}$$

with $\alpha = 0.06$.

Algorithm 1: Harris Corner Detector

Compute the horizontal and vertical derivatives I_x and I_y of the image by convolving the original image with a Sobel filter;
Compute the three images corresponding to the outer products of these gradients. (The matrix A is symmetric, so only three entries are needed.);
Convolve each of these images with a larger Gaussian.;
Compute a scalar interest measure using the formulas (Equation 2) discussed above.;
Find local maxima above a certain threshold and report them as detected feature point locations.;

To implement the Harris corner detector, you will have to fill out the following methods in `part1_harris_corner.py`:

- `compute_image_gradients()`: Computes image gradients using the Sobel filter.
- `get_gaussian_kernel_2D_pytorch()`: Creates a 2D Gaussian kernel (this is essentially the same as your Gaussian kernel method from project 1).
- `second_moments()`: Computes the second moments of the input image. You will need to use your `get_gaussian_kernel_2D_pytorch()` method.
- `compute_harris_response_map()`: Gets the raw corner responses over the entire image (the previously implemented methods may be helpful).
- `maxpool_numpy()`: Performs the maxpooling operation using just NumPy. This manual implementation will help you understand what's happening in the next step.
- `nms_maxpool_pytorch()`: Performs non-maximum suppression using max-pooling. You can use PyTorch max-pooling operations for this.
- `remove_border_vals()`: Removes values close to the border that we can't create a useful SIFT window around.
- `get_harris_interest_points()`: Gets interests points from the entire image (the previously implemented methods may be helpful).

The starter code gives some additional suggestions. You do not need to worry about scale invariance or keypoint orientation estimation for your baseline Harris corner detector. The original paper by Chris Harris and Mike Stephens describing their corner detector can be found here.

2 Part 2: Local feature descriptors (`part2_patch_descriptor.py`)

To get your matching pipeline working quickly, you will implement a bare-bones feature descriptor in `part2_patch_descriptor.py` using normalized, grayscale image intensity patches as your local feature. See Szeliski 7.1.2 for more details when coding `compute_normalized_patch_descriptors()`

Choose the top-left option of the 4 possible choices for center of a square window, as shown in Figure 2.

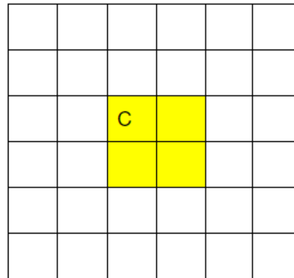

Figure 2: For this example of a 6×6 window, the yellow cells could all be considered the center. Please choose the top left (marked "C") as the center throughout this project.

3 Part 3: Feature matching (`part3_feature_matching.py`)

You will implement the "ratio test" (also known as the "nearest neighbor distance ratio test") method of matching local features as described in the lecture materials and Szeliski 7.1.3 (page 421). See equation 7.18 in particular. The potential matches that pass the ratio test the easiest should have a greater tendency to be correct matches – think about *why* this is. In `part3_feature_matching.py`, you will have to code `compute_feature_distances()` to get pairwise feature distances, and `match_features_ratio_test()` to perform the ratio test to get matches from a pair of feature lists.

4 Part 4: SIFT Descriptor (`part4_sift_descriptor.py`)

You will implement a SIFT-like local feature as described in the lecture materials and Szeliski 7.1.2. We'll use a simple one-line modification ("Square-Root SIFT") from a 2012 CVPR paper (linked here) to get a free boost in performance. See the comments in the file `part4_sift_descriptor.py` for more details.

Regarding Histograms SIFT relies upon histograms. An unweighted 1D histogram with 3 bins could have bin edges of $[0, 2, 4, 6]$. If $x = [0.0, 0.1, 2.5, 5.8, 5.9]$, and the bins are defined over half-open intervals $[e_{left}, e_{right})$ with edges e, then the histogram $h = [2, 1, 2]$.

A weighted 1D histogram with the same 3 bins and bin edges has each item weighted by some value. For example, for an array $x = [0.0, 0.1, 2.5, 5.8, 5.9]$, with weights $w = [2, 3, 1, 0, 0]$, and the same bin edges ($[0, 2, 4, 6]$), $h_w = [5, 1, 0]$. In SIFT, the histogram weight at a pixel is the magnitude of the image gradient at that pixel.

In `part4_sift_descriptor.py`, you will have to implement the following:

- `get_magnitudes_and_orientations()`: Retrieves gradient magnitudes and orientations of the image.
- `get_gradient_histogram_vec_from_patch()`: Retrieves a feature consisting of concatenated histograms.

- `get_feat_vec()`: Gets the adjusted feature from a single point.
- `get_SIFT_descriptors()`: Gets all feature vectors corresponding to our interest points from an image.

5 Writeup

For this project (and all other projects), you must do a project report using the template slides provided to you. Do ***not*** change the order of the slides or remove any slides, as this will affect the grading process on Gradescope and you will be deducted points. In the report you will describe your algorithm and any decisions you made to write your algorithm a particular way. Then you will show and discuss the results of your algorithm. The template slides provide guidance for what you should include in your report. A good writeup doesn't just show results – it tries to draw some conclusions from the experiments. You must convert the slide deck into a PDF for your submission.

If you choose to do anything extra, add slides *after the slides given in the template deck* to describe your implementation, results, and analysis. Adding slides in between the report template will cause issues with Gradescope, and you will be deducted points. You will not receive full credit for your extra credit implementations if they are not described adequately in your writeup. In addition, when turning in the PDF writeup to gradescope, please match the pages of the writeup to the appropriate sections of the rubric.

Using the starter code (`proj2.ipynb`)

The top-level iPython notebook, `proj2.ipynb`, provided in the starter code includes file handling, visualization, and evaluation functions for you, as well as calls to placeholder versions of the three functions listed above.

For the Notre Dame image pair there is a ground truth evaluation in the starter code as well. `evaluate_correspondence()` will classify each match as correct or incorrect based on hand-provided matches (see `show_ground_truth_corr()` for details). The starter code also contains ground truth correspondences for two other image pairs (Mount Rushmore and Episcopal Gaudi). You can test on those images by uncommenting the appropriate lines in `proj2.ipynb`.

As you implement your feature matching pipeline, you should see your performance according to `evaluate_correspondence()` increase. Hopefully you find this useful, but don't *overfit* to the initial Notre Dame image pair, which is relatively easy. The baseline algorithm suggested here and in the starter code will give you full credit and work faily well on these Notre Dame images, but additional image pairs provided in `extra_data.zip` are more difficult. They might exhibit more viewpoint, scale, and illumination variation.

Potentially useful NumPy and Pytorch functions

From Numpy: `np.argsort()`, `np.arctan2()`, `np.concatenate()`, `np.fliplr()`, `np.flipud()`, `np.histogram()`, `np.hypot()`, `np.linalg.norm()`, `np.linspace()`, `np.newaxis`, `np.reshape()`, `np.sort()`.

From Pytorch: `torch.argsort()`, `torch.arange()`, `torch.from_numpy()`, `torch.median()`, `torch.nn.functional.conv2d()`, `torch.nn.Conv2d()`, `torch.nn.MaxPool2d()`, `torch.nn.Parameter`, `torch.stack()`.

For the optional, extra-credit vectorized SIFT implementation, you might find `torch.meshgrid`, `torch.norm`, `torch.cos`, `torch.sin`.

We want you to build off of your Project 1 expertise. Please use `torch.nn.Conv2d` or `torch.nn.functional.conv2d` instead of convolution/cross-correlation functions from other libraries (e.g., `cv.filter2D()`, `scipy.signal.convolve()`).

Forbidden functions

(You can use these OpenCV, Sci-kit Image, and SciPy functions for testing, but not in your final code). `cv2.getGaussianKernel()`,`np.gradient()`, `cv2.Sobel()`, `cv2.SIFT()`, `cv2.SURF()`, `cv2.BFMatcher()`, `cv2.BFMatcher().match()`, `cv2.BFMatcher().knnMatch()`, `cv2.FlannBasedMatcher().knnMatch()`, `cv2.HOGDescriptor()`, `cv2.cornerHarris()`, `cv2.FastFeatureDetector()`, `cv2.ORB()`, `skimage.feature`, `skimage.feature.hog()`, `skimage.feature.daisy`, `skimage.feature.corner_harris()`, `skimage.feature.corner_shi_tomasi()`, `skimage.feature.match_descriptors()`, `skimage.feature.ORB()`, `cv.filter2D()`, `scipy.signal.convolve()`.

We haven't enumerated all possible forbidden functions here, but using anyone else's code that performs interest point detection, feature computation, or feature matching for you is forbidden.

Tips, tricks, and common problems

- Make sure you're not swapping x and y coordinates at some point. If your interest points aren't showing up where you expect, or if you're getting out of bound errors, you might be swapping x and y coordinates. Remember, images expressed as NumPy arrays are accessed `image[y, x]`.
- Make sure your features aren't somehow degenerate. you can visualize features with `plt.imshow(image1_features)`, although you may need to normalize them first. If the features are mostly zero or mostly identical, you may have made a mistake.

Bells & whistles (extra points) / Extra Credit

Implementation of bells & whistles can increase your grade on this project by up to 10 points (potentially over 100). The max score for all students is 110.

For all extra credit, be sure to include quantitative analysis showing the impact of the particular method you've implemented. Each item is "up to" some amount of points because trivial implementations may not be worthy of full extra credit.

Local feature description

- up to 3 pts: The simplest thing to do is to experiment with the numerous SIFT parameters: How big should each feature be? How many local cells should it have? How many orientations should each histogram have? Different normalization schemes can have a significant effect as well. Don't get lost in parameter tuning though.
- up to 10 pts: Implement a vectorized version of SIFT that runs in under 5 seconds, with at least 80% accuracy on the Notre Dame image pair.

Rubric

- +25 pts: Implementation of Harris corner detector in `part1_harris_corner.py`
- +10 pts: Implementation of patch descriptor `part2_patch_descriptor.py`
- +10 pts: Implementation of "ratio test" matching in `part3_feature_matching.py`
- +35 pts: Implementation of SIFT-like local features in `part4_sift_descriptor.py`
- +20 pts: Report
- -5*n pts: Lose 5 points for every time you do not follow the instructions for the hand-in format

Submission format

This is very important as you will lose 5 points for every time you do not follow the instructions. You will submit two items to Gradescope:

1. `<your_gt_username>.zip` containing:
   (a) `proj2_code/` - directory containing all your code for this assignment
   (b) `additional_data/` - (optional) if you use any data other than the images we provide, please include them here
2. `<your_gt_usernamme>_proj2.pdf` - your report

Do ***not*** install any additional packages inside the conda environment. The TAs will use the same environment as defined in the config files we provide you, so anything that's not in there by default will probably cause your code to break during grading. Do ***not*** use absolute paths in your code or your code will break. Use relative paths like the starter code already does. Failure to follow any of these instructions will lead to point deductions. Create the zip file using `python zip_submission.py --gt_username <your_gt_username>` (it will zip up the appropriate directories/files for you!) and hand it in with your report PDF through Gradescope (please remember to mark which parts of your report correspond to each part of the rubric).

Credits

Assignment developed by James Hays, Cusuh Ham, John Lambert, Vijay Upadhya, and Samarth Brahmbhatt.