

Lucas-Kanade Tracker

Titus Cieslewski

Contents

1 Preliminaries	1
1.1 Outline of the exercise	1
1.2 Provided code	2
1.3 Conventions	2
2 Part 1: Warping images	2
3 Part 2: Recovering a simple warp with brute force	3
4 Part 3: Recovering the warp with KLT	4
5 Part 4: Applying KLT to KITTI	6
6 Part 5: Outlier rejection with the bidirectional error	6

In this exercise, we will replace the crude matching method from exercise 3 with a Lucas-Kanade Tracker.

1 Preliminaries

1.1 Outline of the exercise

As seen in today's lecture, a Lukas-Kanade Tracker (KLT¹) can be used to track the position of a point accross frames. In this exercise, we will implement KLT from scratch. We will start with a well-defined problem where we warp an input image ourselves and see whether KLT can recover the warp. Then, we will apply our KLT to the KITTI sequence and see how it improves on the matching from exercise 3.

¹for some reason the names are inverted in the acronym



Figure 1: Tracking 50 keypoints with the KLT implementation of this exercise. See <https://youtu.be/iaRPafeG9zw> for a video. In this exercise, we apply KLT to downsampled images, since implementing the coarse-fine estimation is outside of the scope of this exercise.




Figure 2: Results for the different image warps in part 1.

1.2 Provided code

As usual, `main.m` contains one section for each part of the exercise which you can run individually with `Ctrl+Enter`. Functions to be implemented are provided as documented stubs and are called in `main.m`.

1.3 Conventions

We depart from the notation used in class only slightly, to better distinguish between the **image** in which the template is defined (reference image) and the **coefficients** of the patch of the template. We call the reference image I_R and the patch coefficients

$$T = \{\vec{x} | \vec{x} = \vec{x}_T + \Delta\vec{x} \quad \forall \vec{x} \in [-r_T, r_T] \times [-r_T, r_T] \subset \mathbb{Z} \times \mathbb{Z}\} \quad (1)$$

where we call \vec{x}_T the patch center and r_T the patch radius, consistently with previous exercises.

Throughout the exercise, you will need to pay attention to **indexing**! Confusion can ensue because we formulate most things as (x, y) , but matrix, and in particular image elements are accessed with (row, column), which corresponds to (y, x) !

2 Part 1: Warping images

Given an image $I(x, y)$, we say $I(W(x, y, \vec{p}))$ is the image I warped by the warp W which is itself parametrized with the parameters \vec{p} . We define the warp with a 2×3 matrix \mathbf{W} such that

$$W(x, y, \vec{p}) = \mathbf{W} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} p_1 & p_3 & p_5 \\ p_2 & p_4 & p_6 \end{bmatrix}. \quad (2)$$

As seen in class, this is called an affine warp, and generalizes translation, rigid and similarity warps, which are each a special case of the affine warp. In order to get a better understanding of KLT and facilitate debugging, we will manually warp an image with \mathbf{W} . Unfortunately, it is not very intuitive to directly manipulate the coefficients of \mathbf{W} . Let us instead provide an interface which allows us to get a similarity warp parametrized by the translation $(\Delta x, \Delta y)$, a rotation angle α and the scale change λ . Write the function `getSimWarp` which returns \mathbf{W} such that

$$W(x, y, \vec{p}) = \lambda \left(\begin{bmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \right). \quad (3)$$

Then, write the function `warpImage` which warps the input image given \mathbf{W} . Note that $W(x, y, \vec{p})$ is likely not to return integer coefficients. You can for now return $I(\lfloor W(x, y, \vec{p}) \rfloor)$, later this will need to be changed to bilinear interpolation to increase robustness of KLT. Also, note that $W(x, y, \vec{p})$ could land outside the image dimensions. In that case just return an intensity of 0. If all goes well, you should obtain the results as in Figure 2.




Figure 3: Template and SSDS for Part 2.

3 Part 2: Recovering a simple warp with brute force

An essential part of the KLT implementation will be extracting warped patches:

$$I(W(T, \vec{p})) := \{I(W(x, y, \vec{p})) \mid \forall (x, y) \in T\}, \quad (4)$$

see (1) for the definition of the patch T . However, there is a slight issue with the current way in which we apply warps. As you can see in Figure 2, rotation and scaling are applied in the top left corner. This makes sense, since that is the origin. But it is also problematic, since for the warping of any image patch not in the top left of the image, rotation and translation are now coupled. This is tedious to think about, but also problematic for the optimization process. In order to avoid this problem, we re-define $I(W(T, \vec{p}))$ such that the origin of the warp is now at the center of the image patch \vec{x}_T :

$$I(W(T, \vec{p})) := \{I(\vec{x}_T + W(\Delta\vec{x}, \vec{p})), \quad \Delta\vec{x} \in [-r_T, r_T]^2 \subset \mathbb{Z}^2\}, \quad (5)$$

where T, \vec{x}_T and $\Delta\vec{x}$ express the image patch, see 1. Implement the function `getWarpedPatch`, which returns $I(W(T, \vec{p}))$. As in the previous part, you can for now return $I(\lfloor W(T, \vec{p}) \rfloor)$, but Parts 4 and 5 will not really work well with that. For those parts, you should use bilinear interpolation. You can copy bilinear interpolation from the reference solution of `undistortImage` in exercise 1.

We can now implement a first, simple, brute force tracker, as seen in the first part of the lecture. The goal of tracking is to find a warp \mathbf{W} which minimizes the sum of square difference (SSD) E between the template and a warped patch from the image in which the template is supposed to be tracked:

$$E(\mathbf{W}) = \sum (I(W(T, \mathbf{W})) - I_R(T))^2 \quad (6)$$

Because the warp is also performed on the patch domain T , we can simply extract how the center point \vec{x}_T has moved from \mathbf{W} : $\Delta\vec{x}_T = [p_5 \ p_6]^T$. Implement `trackBruteForce`, which recovers **translation-only** warp by trying out different values for (p_5, p_6) and selects the one with the lowest such SSD:

$$\operatorname{argmin} E(\mathbf{W}'(p_5, p_6)), \quad \mathbf{W}' = \begin{bmatrix} 1 & 0 & p_5 \\ 0 & 1 & p_6 \end{bmatrix}, \quad (7)$$

where the range of (p_5, p_6) to try is itself defined as a patch D with parameters $(x_D, y_D) = (x_T, y_T)$ but a different radius r_D . For the given test cases, the plots should look like in Figure 3. `main.m` gives you feedback on what the recovered (p_5, p_6) should be.

4 Part 3: Recovering the warp with KLT

As discussed in the lecture, brute force tracking works well, but has the disadvantage that it has a large space of possible \mathbf{W} to explore. In the previous part, this space was mitigated to two dimensions by only using two parameters for \mathbf{W} . But the size of the space grows exponentially with the parameters, and a general affine warp has six parameters!

Thankfully, it turns out that $E(\mathbf{W})$ is mostly smooth and locally convex around the optimal solution $E(\mathbf{W}^*)$. This is illustrated for (p_5, p_6) aka $(\Delta x, \Delta y)$ on the right side of Figure 3. This local convexity suggests that we can use gradient descent to find \mathbf{W}^* , assuming that we start with an initial guess that is sufficiently close to \mathbf{W}^* . KLT uses the Gauss-Newton method for gradient descent. As seen in class, this is first formulated as finding an increment $\Delta \vec{p}$ to that initial guess \vec{p} that minimizes the error E :

$$E = \sum (I(W(T, \vec{p} + \Delta \vec{p})) - I_R(T))^2 \quad (8)$$

$$E(\Delta \vec{p}) = \min E \Leftrightarrow \frac{\partial}{\partial \Delta \vec{p}} \sum (I(W(T, \vec{p} + \Delta \vec{p})) - I_R(T))^2 = 0 \quad (9)$$

If we wanted to solve this directly, we would need to find a differentiable expression for $I(W(T, \vec{p} + \Delta \vec{p}))$. You are welcome to try, but you will find that this requires a pixel-wise case distinction. This would require much more computation than just finding \vec{p} with the above brute force method. The Gauss-Newton algorithm avoids this problem by linearizing around \vec{p} , that is, making the assumption $I(W(T, \vec{p} + \Delta \vec{p})) \sim I(W(T, \vec{p})) + \frac{\partial I}{\partial \vec{p}} \Delta \vec{p}$ (first-order Taylor approximation). Using this approximation will result in an error, but solving the approximation iteratively should converge to the local minimum². With this approximation, (9) becomes

$$\frac{\partial}{\partial \Delta \vec{p}} \sum (I(W(T, \vec{p})) + \frac{\partial I}{\partial \vec{p}} \Delta \vec{p} - I_R(T))^2 = 0 \quad (10)$$

which we can now derive with respect to $\Delta \vec{p}$. For simplicity of expression, let us use vector calculus. We define the column vectors $\vec{e}, \vec{i}, \vec{i}_R$ which are the vectorizations of the expression inside the square root in (10), $I(W(T, \vec{p}))$ and $I_R(T)$, respectively. Each coefficient of these vectors then corresponds to a pixel of the patch T and we have:

$$\vec{e} = \vec{i} + \frac{\partial \vec{i}}{\partial \vec{p}} \Delta \vec{p} - \vec{i}_R : \quad (11)$$

Recall that a central component of vector calculus is the Jacobian matrix $\frac{\partial \vec{f}}{\partial \vec{x}}$, which expresses the derivative of a function $\vec{f}(\vec{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ as follows:

$$\frac{\partial \vec{f}}{\partial \vec{x}} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_n} \end{bmatrix} \quad (12)$$

Note that $\frac{\partial \vec{f}}{\partial \vec{x}}$ is $(m \times n)$. It is conveniently defined such that $f(\vec{x}) = A\vec{x} \Rightarrow \frac{\partial \vec{f}}{\partial \vec{x}} = A$.

So, with (11), we can rewrite (10) as:

$$\frac{\partial}{\partial \Delta \vec{p}} \vec{e}^T \vec{e} = \vec{0}. \quad (13)$$

Let us derive with respect to $\Delta \vec{p}$. Remember that \vec{e} depends on $\Delta \vec{p}$, so we need to apply the chain rule:

$$\frac{\partial}{\partial \Delta \vec{p}} \vec{e}^T \vec{e} = \frac{\partial}{\partial \vec{e}} \vec{e}^T \vec{e} \cdot \frac{\partial \vec{e}}{\partial \Delta \vec{p}} \quad (14)$$

²See Newton's method for finding the root (zero-crossing) of one-dimensional functions for some intuitive understanding.

It is simple to show that $\frac{\partial}{\partial \vec{e}} \vec{e}^T \vec{e} = 2\vec{e}^T$. You should be able to show that with pen and paper by writing out the coefficients of $\vec{e}^T \vec{e}$ and then writing out the Jacobian $\frac{\partial}{\partial \vec{e}}$ for these coefficients. As for $\frac{\partial \vec{e}}{\partial \Delta \vec{p}}$, it is simply $\frac{\partial I}{\partial \vec{p}}$, as can be seen from (11). So we now have:

$$2\vec{e}^T \cdot \frac{\partial \vec{i}}{\partial \vec{p}} = \vec{0} \quad (15)$$

Dividing by two and applying $(AB)^T = B^T A^T$ for later convenience, we get:

$$\frac{\partial \vec{i}^T}{\partial \vec{p}} \cdot \vec{e} = \frac{\partial \vec{i}^T}{\partial \vec{p}} (\vec{i} + \frac{\partial \vec{i}}{\partial \vec{p}} \Delta \vec{p} - \vec{i}_R) = \vec{0} \quad (16)$$

and can now proceed to solve for $\Delta \vec{p}$, by taking it out from the product, moving the other terms to the right side and eliminating the factor $\Delta \vec{p}$ is multiplied with:

$$\frac{\partial \vec{i}^T}{\partial \vec{p}} (\vec{i} + \frac{\partial \vec{i}}{\partial \vec{p}} \Delta \vec{p} - \vec{i}_R) = \underbrace{\frac{\partial \vec{i}^T}{\partial \vec{p}} \frac{\partial \vec{i}}{\partial \vec{p}}}_{H} \Delta \vec{p} + \frac{\partial \vec{i}^T}{\partial \vec{p}} (\vec{i} - \vec{i}_R) = \vec{0} \Leftrightarrow \quad (17)$$

$$H \Delta \vec{p} = \frac{\partial \vec{i}^T}{\partial \vec{p}} (\vec{i}_R - \vec{i}) \Leftrightarrow \Delta \vec{p} = H^{-1} \frac{\partial \vec{i}^T}{\partial \vec{p}} (\vec{i}_R - \vec{i}), \quad (18)$$

where we have introduced H , the so-called Hessian. With this derivation, you should be able to write the function `trackKLT`, which iteratively updates $\vec{p} \leftarrow \vec{p} + \Delta \vec{p}$ with the update rule we just derived. Start with a \vec{p} which corresponds to an identity warp. Note that $\frac{\partial \vec{i}}{\partial \vec{p}}$ can be further decomposed as seen in the class. Because $\vec{i}(\vec{p}) = \vec{i}(\vec{w})$, $\vec{w} = \mathbf{W}\vec{x}$, you can use the chain rule: $\frac{\partial \vec{i}}{\partial \vec{p}} = \frac{\partial \vec{i}}{\partial \vec{w}} \cdot \frac{\partial \vec{w}}{\partial \vec{p}} = \begin{bmatrix} \vec{i}_x & \vec{i}_y \end{bmatrix} \frac{\partial}{\partial \vec{p}} W(x, y, \vec{p}) = \begin{bmatrix} \vec{i}_x & \vec{i}_y \end{bmatrix} \begin{bmatrix} \Delta x & 0 & \Delta y & 0 & 1 & 0 \\ 0 & \Delta x & 0 & \Delta y & 0 & 1 \end{bmatrix}$ where $\Delta x, \Delta y$ are the pixel coefficients relative to \vec{x}_T , see (5). Here, \vec{i}_x, \vec{i}_y are nothing but the x- and y-gradients of the **warped** image patch, vectorized. This all sounds good until you try to implement it and realize that you can't simply implement it as a matrix multiplication, because $\Delta x, \Delta y$ have different values for different pixels, that is, different rows of $\begin{bmatrix} \vec{i}_x & \vec{i}_y \end{bmatrix}$. So, to construct $\frac{\partial \vec{i}}{\partial \vec{p}}$, you will need to multiply each row of $\begin{bmatrix} \vec{i}_x & \vec{i}_y \end{bmatrix}$ with a numerically different $\frac{\partial \vec{w}}{\partial \vec{p}}$. Hints:

- Use `getWarpedPatch` to get I and I_R , then obtain \vec{i} and \vec{i}_R by vectorizing the matrices.
- Similarly, it is convenient to obtain the patch gradients by convolving I first, and vectorizing the patch later. For convenience, you can get the image gradients by convolving with $\begin{bmatrix} 1 & 0 & -1 \end{bmatrix}$ rather than with $\begin{bmatrix} 1 & -1 \end{bmatrix}$. Note that this convolution, if applied validly, will reduce the patch size by one pixel on each side. To make sure that the gradient is defined on the full patch, use as input to the convolutions an accordingly larger patch.
- A good first step when debugging is to check the whether the shapes of your intermediate products make sense. For example, in (18) the shapes of matrices H^{-1} , $\frac{\partial \vec{i}^T}{\partial \vec{p}}$, $(\vec{i}_R - \vec{i})$ should be, respectively, (6×6) , $(6 \times |T|)$, $(|T| \times 1)$, where $|T|$ is the amount of pixels in the patch.
- You should be able to write down convenient intermediate products (try to avoid redundant calculations!) from the above derivation. If you struggle, consult the lecture notes.
- For debugging and visualization, we recommend to plot $I_R, I, (I_R - I)$, the warped image gradients I_x, I_y and the steepest descent images $\nabla I \frac{\partial W}{\partial \vec{p}}$, which you can obtain by reshaping each column of $\frac{\partial \vec{i}}{\partial \vec{p}}$ into an image patch. See Figure 4 and <https://youtu.be/3g-zEHoVwvM> for how these look like in our implementation.




Figure 4: Visualization of KLT gradient descent for the very first iteration. See <https://youtu.be/3gzEHoVwvM> for the video.

5 Part 4: Applying KLT to KITTI

You should now be able to run the next section, in which we take the keypoint tracker and apply it to the KITTI dataset. This will work much better if you have implemented bilinear interpolation in `getWarpedPatch`. Note that we have downsampled the images by a factor of 4. This ensures that a sufficient amount of them are tracked; with a higher resolution, the motion of too many points is outside of the convergence basin of KLT. As you should see, some points are still not correctly tracked. Next, we will implement a simple method to discard points which are not correctly tracked.

6 Part 5: Outlier rejection with the bidirectional error

The bidirectional error check is a simple test to verify that a point is consistently tracked. Let us define the result of tracking a point \vec{x}_T from image I_i to image I_{i+1} with KLT as

$$\Psi_i^{i+1}(\vec{x}_T) = \vec{x}_T + \begin{bmatrix} p_5 \\ p_6 \end{bmatrix}. \quad (19)$$

Then, the bidirectional error test verifies that when we try to track the point back to I_i , we are close to \vec{x}_T again:

$$\Psi_{i+1}^i(\Psi_i^{i+1}(\vec{x}_T)) - \vec{x}_T < \lambda \quad (20)$$

Implement `trackKLTRobustly` which encapsulates $\Psi_i^{i+1}(\vec{x}_T)$ and also indicates whether the tracking has passed the bidirectional error test. Then, when you run the corresponding section of `main.m`, you should get something like in Figure 1 and <https://youtu.be/iaRPafeG9zw>.