LAB ASSIGNMENT 6

ETH ZURICH, COMPUTER SCIENCE DEPARTMENT

Computer vision

Report

Condensation tracker

Author Emilia Szymańska, 22-945-547

December 2022

The objective of this lab assignment was to implement a CONDENSATION tracker based on color histograms and to run experiments to test the influence of the parameters for different types of videos on the performance.

1 CONDENSATION Tracker Based On Color Histograms

1.1 Color histogram

The first step was to calculate the normalized histogram of RGB colors within the defined bounding box. I slice the frame in accordance to the minimum and maximum coordinates provided as our bounding box description. Then I separate the RGB channels and for each of them I calculate the color histogram with the provided number of bins. Lastly, I normalize the histograms by diving them by their sums and create an output array containing histograms for three channels.

The code implementation of this part is presented below:

```
1 def color histogram(xmin, ymin, xmax, ymax, frame, hist bin):
2
3
       sliced img = frame[ymin:ymax,xmin:xmax]
4
                      = sliced_img[:,:,0].flatten()
5
       red channel
6
       green_channel = sliced_img[:,:,1].flatten()
7
       blue_channel = sliced_img[:,:,2].flatten()
8
                   \_ = np.histogram(red_channel, bins=hist bin)
9
       hist red,
       hist_green , _ = np.histogram(green_channel, bins=hist_bin)
hist_blue , _ = np.histogram(blue_channel, bins=hist_bin)
10
11
12
       hist red
13
                 = hist_red/np.sum(hist_red)
       hist green = hist green/np.sum(hist green)
14
15
       hist_blue = hist_blue/np.sum(hist_blue)
16
       hist = np.array([hist_red, hist_green, hist_blue])
17
18
19
       return hist
```

1.2 Deriving A matrix

The state we consider is in form:

$$s = \{x, y, \dot{x}, \dot{y}\}\tag{1}$$

where x, y represent the center of the bounding box and \dot{x}, \dot{y} are the velocities in x and y directions. Therefore, for no motion prediction model we do not want to change the values of the input vector, so A should be an identity matrix.

When we consider the constant velocity motion model, we want to add \dot{x}, \dot{y} to x, y respectively. Therefore the matrix A is of form:

$$A = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
(2)

1.3 Propagation

This function propagates the particles given the system prediction model (matrix A) and the system model noise parameters in accordance to equation s' = As + w. First I define A matrix for both cases as well as the noise vector which is a normal distribution with the given sigma (varying for position and velocity values). New particles are computed by multiplying the matrix A with old particles and adding the noise vector. Finally, I make sure that the new particles lie within image dimensions.

The code implementation of this part is presented below:

```
1 def propagate (particles, frame height, frame width, params):
2
       if params ['model'] == 0: \# no motion
3
4
           A = np.identity(2)
           w = np.transpose(np.random.normal(0, params['sigma position'], size=
5
               particles.shape))
       else:
                                  \# motion with const vel
6
           A = np.array( [ [ 1, 0, 1, 0 ], [ 0, 1, 0, 1 ], [ 0, 0, 1, 0 ], [ 0, 0, 0, 1 ] ])
7
           w = np.transpose(np.random.normal(0, [params['sigma_position'], params['])
8
               sigma_position'], params['sigma_velocity'], params['sigma_velocity']],
                size=particles.shape))
9
       new_particles = np.dot(A, np.transpose(particles)) + w
10
11
       new_particles = np.transpose(new_particles)
12
13
       new_particles = np.where(new_particles < 0, 0, new_particles)
14
       for part in new_particles:
15
16
           if part [0] > \text{frame width} -1:
               part[0] = frame_width-1
17
           if part[1] > frame_height-1:
18
19
               part[1] = frame height - 1
20
       return new_particles
21
```

1.4 Observation

In the observe() function, for each particle I calculate the bounding box corners' minimum and maximum coordinates (and make sure they lie within the frame. Then I calculate the color histogram for this bounding box and later I compute the weight in accordance to the formula:

$$w = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{\chi^2 (CH_s, CH_{target})^2}{2\sigma^2}} \tag{3}$$

I store all the weights in a list and once I compute them for all the particles, I create a numpy array from them and normalize them.

The code implementation of this part is presented below:

```
1 def observe(particles, frame, bbox height, bbox width, params hist bin, hist,
       params sigma observe):
2
       height, width, _ = frame.shape
3
4
       norm term = 1/(\text{math.sqrt}(2*\text{math.pi})*\text{params sigma observe})
       weights = []
5
6
7
       for part in particles:
8
           x_{max} = int(part[0]+bbox_width/2)
           x_{min} = int(part[0]-bbox_width/2)
9
10
           if x max > width -1:
                x max = width - 1
11
12
           if x \min < 0:
13
                x \min = 0
14
           y max = int(part[1]+bbox height//2)
15
           y_{min} = int(part[1]-bbox_height//2)
16
17
           if y_{max} > height -1:
               y max = height - 1
18
19
           if y_{min} < 0:
                y_{min} = 0
20
21
           hist observed = color histogram(x min, y min, x max, y max, frame,
22
               params hist bin)
```

1.5 Estimation

In this step, I compute the estimate of the mean state, which is a sum of the product of particles and their weights.

The code implementation of this part is presented below:

```
1 def estimate(particles, particles_w):
2 return np.sum(particles_w*particles, axis=0)
```

1.6 Resampling

Finally, in accordance to the particles' weights, I randomly select new particles. The corresponding weights need to be normalized before being returned.

The code implementation of this part is presented below:

```
1 ...
2 def resample (particles, particles w):
      chosen ids = np.random.choice(np.arange(particles.shape[0]), particles.shape
3
          [0], p=particles_w.flatten())
4
                        = np.take(particles, chosen_ids, axis=0)
5
      chosen_particles
6
      chosen_particles_w = np.take(particles_w, chosen_ids, axis=0)
7
      chosen_particles_w = chosen_particles_w/np.sum(chosen_particles_w)
8
9
      return \ chosen\_particles \ , \ chosen\_particles\_w
10 ...
```

2 Experiments

2.1 Video 1

The first video runs with the tracker as presented in Fig. 1 for no motion model and Fig. 2 for constant velocity model. It can be seen that the tracker manages to follow the top part of the hand, which was selected for tracking, however it was more noisy for the constant velocity model. The blue trajectory corresponds to a priori mean state and the red trajectory to a posteriori mean state.



Figure 1: Result of running the tracker with no motion model on the video nr 1 (tracking the hand).



Figure 2: Result of running the tracker with constant velocity model on the video nr 1 (tracking the hand).

2.2Video 2

The second video's objective is to track a hand, which at some point is occluded by another object. This time I was supposed to vary the parameters to obtain the final tracking (Fig. 3). I run the video for different σ values and modes of operation (Fig. 5, 4). The tracking had better results for constant velocity model – for no motion model the correct tracking was usually stopped when encountering the occluding obstacle. Increasing $\sigma_{position}$ resulted in highly "scattering" behaviour, but decreasing it made the tracker prefer states without almost any movement. With high $\sigma_{velocity}$, the tracker performed big steps between frames, the opposite happening for low $\sigma_{velocity}$ values. As for measurement noise, using a high $\sigma_{observe}$ resulted in a priori and a posteriori trajectories to be almost exactly the same, therefore not following the object. Significantly decreasing $\sigma_{observe}$, however, will also not be a solution - the trajectories are also chaotic (although more accurate than in the other case).



Figure 3: Tracking the hand on the video nr 2.



(b) No motion

(c) Constant velocity



Figure 4: Comparison of no motion and constant velocity model for the tracker on video nr 2.



Figure 5: Comparison of σ magnitudes for velocity, position and observation for video nr 2.

2.3 Video 3

In my case, plugging the parameters from the video nr 2 to the video nr 3 did not allow for tracking the ball (Fig. 6) – the tracker got "stuck" after the ball changed its direction of movement. I needed to tune the parameters to make it possible to properly follow the ball's movement (Fig. 7). The comparison for $\sigma_{position}$ values for this video yields similar results to the $\sigma_{position}$ varying of video nr 2 – it affects the "willingness" to change the position of the center of the bounding box. It is however slightly harder to spot the difference here when varying $\sigma_{velocity}$ – we can see that with a lower value it is more likely to be "stuck" in the very corner of the image. High $\sigma_{observe}$ value again makes the a priori and a posteriori trajectories to be equal and they do not move fully to the direction of the ball movement (they however managed in this case to partially "find" the ball again when it started coming back to the initial position), while decreasing it results in more noisy behaviour of both trajectories. The comparison of σ values and their effects on the videos is preented in Fig. 8. Considering changing the model from no motion to constant velocity one, there is a significant difference in the result. The good trajectory in Fig. 7 was based on the no motion model and constant velocity one resulted in similar behaviour to Fig. 6 – the tracker stopped tracking when the ball hit the wall and changed the movement.



Figure 6: The effect of using the parameters tuned for video 2 on video 3.



Figure 7: Tracking the ball on video 3.



(e) High $\sigma_{observe}$

(f) Low $\sigma_{observe}$

Figure 8: Comparison of σ magnitudes for velocity, position and observation for video nr 3.

2.4 General observations

In all the videos, there were parameters which had similar results when being changed. Having fewer bins in histogram color model made the tracking more general, therefore it was less accurate (we made the disctinction on fewer features). If we increase this number, we will obtain better results, however at some point the tracker can become too specific when looking for similarities and it may not be accurate anymore. Allowing the appearance model to be updated after each iteration adjusts it to the current surrounding, which in a varying environment might be useful (in our case we mostly had plain wall throughout the video, but if on the way the background is significantly different than the starting one, we might lose the correct tracking). The number of particles affects the choice of the center of the tracking bounding box – again, usually more is better for tracking the object, however too many may result in too specific search and taking into consideration data which is not relevant. If we do not have enough particles, the result will be highly noisy because of the lack of informative data.