CS 91R: The Computational Image

Assignment 8: Foreground-Background Segmentation

DUE April 1st at 11:59 PM

In this assignment we will compare techniques for image segmentation, in particular, modeling the background in a video sequence. Each method models the background with increasing sophistication. Once the background model is computed, a pixel can be classified as background or foreground using a simple threshold decision rule. For this lab, implement and compare the following methods using grayscale or RGB/HSV images.

1. **Baseline:** Take a static image (image 0) and subtract that from subsequent frames. Then decide a pixel x is foreground only if the intensity difference is more than D:

$$\|x - x_0\| > D$$

- **Note:** x refers to a pixel: either grayscale intensity, RGB, or HSV, not a x-coordinate.

2. **Average:** Compute an average value for each pixel's intensity over time. A running average can be computed using the following efficient update rule.

$$\hat{\mu}_t = \frac{1}{t}\sum_{i=1}^{t} x_i, \qquad \hat{\mu}_t = \frac{t-1}{t}\mu_{t-1} + \frac{1}{t}x = \hat{\mu}_{t-1} + \frac{1}{t}(x - \hat{\mu}_{t-1}) \qquad \|x - \hat{\mu}_t\| > D$$

3. **Moving Average:** Compute a moving-average (again, over time) for each pixel's intensity to model the background (where $0 \leq \alpha \leq 1$).

$$\hat{\mu}_t = (1-\alpha)\hat{\mu}_{t-1} + \alpha x = \hat{\mu}_{t-1} + \alpha(x - \hat{\mu}_{t-1}) \qquad \|x - \hat{\mu}_t\| > D$$

4. **Gaussian:** Assume each pixel's intensity in the background image can be modeled probabilistically using a Gaussian probability distribution. Once we have this model, we can compute the likelihood of a particular pixel value.[1] You can estimate μ and σ^2 from a set of data x_i using the following equations:

$$\hat{\mu}_t = \frac{1}{t}\sum_{i=1}^{t} x_i, \qquad \hat{\sigma_t}^2 = \frac{1}{t-1}\sum_{i=1}^{t}(x_i - \hat{\mu})^2 \qquad p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2 / 2\sigma^2}$$

$$\hat{\mu}_t = \hat{\mu}_{t-1} + \frac{x_t - \hat{\mu}_{t-1}}{t} \qquad S_t = S_{t-1} + (x_t - \hat{\mu}_{t-1})(x_t - \hat{\mu}_t) \qquad \hat{\sigma}_t^2 = \frac{1}{t-1}S_t$$

- This method can used in the same fashion as the previous thresholding functions, but using the *Mahalanobis* distance rather than the absolute difference. The *Mahalanobis* is also known as the z-score in one dimension.

$$r^2 = (\mathbf{x} - \mu)^{\mathsf{T}}\mathbf{\Sigma}^{-1}(\mathbf{x} - \mu) \qquad r = \|\frac{x - \mu_t}{\sigma_t}\| \qquad \|\frac{x - \mu_t}{\sigma_t}\| > D$$

[1] Feel free to use a different pixel model: a) multiple independent univariate Gaussian distributions to describe the RGB or HSV of each pixel b) a multivariate Gaussian distribution c) a Gaussian mixture model.

1 Challenge Problems

1. Try using `img.filter(BLUR)`, `img.filter(ERODE)` or `img.filter(DILATE)` on the difference image to remove some noise.
2. Try the method on HSV (or RGB if you did grayscale).
3. Consider neighboring pixels when deciding if a pixel is in the background.
4. Only update the background model for pixels currently considered background.
5. Come up with some other technique for distinguishing the background from the foreground.

2 Tips

- You can turn off auto-exposure, and control the exposure, on the webcams using the command-line:

```
$ uvcdynctrl -d video0 -s 'Auto Exposure' -- 1
$ uvcdynctrl -d video0 -s 'Exposure Time, Absolute' -- 100
```

3 Learning Objectives

- perform image segmentation on video
- probabilistically model pixels

4 Deliverables

1. Commit the JavaScript `sketch.js` to the repo.
2. Write the reflection (as a markdown document named `reflection.md`) about what you were able to accomplish in this lab. Don't forget the collaboration statement!