

Adaptive Multispectral Difference Weighting

Andrew J. Cosand
Computer Vision and Robotics Research Laboratory
Department of Electrical and Computer Engineering
University of California, San Diego
La Jolla, CA 92093-0407
acosand@ucsd.edu

Abstract

This project attempts to develop a robust and reliable method for fusing hue and intensity difference information for shadow-invariant difference detection. While hue information is standardly used because it is insensitive to shadows, it may fail on pixels which are nearly gray as their hues cannot be determined reliably. A log-scaled distance between each pixel and the gray line in RGB color space is used as a measure of how strongly colored the pixel is. Using information about how the algorithm performs on a training set, weighting functions are generated which fuse hue and intensity difference information as a function of this log-scaled distance. The weighting function currently generated is observably better in some regards than both the hue and saturation differences, but it has yet to combine the better performance characteristics of each without any of the poor performance characteristics. Although the original goal was to fit simple curves to these weighting functions and adaptively adjust the curve parameters, that has not yet been accomplished.

1. Introduction

The goal of background differencing is to use the difference between the current image and a reference image in order to detect interesting objects which appear in the current image, but not in the reference image.

Differencing performance is highly dependant on color space. The HSI color space is often used for differencing in situations where robust detection is required in scenes with shadows, as the hue value is relatively invariant under shadows. Unfortunately, hue determination is unreliable when the image pixel is close to or equal to some shade of gray, because noise or small variations in the image can cause the hue to swing wildly. In theory, a pixel's saturation can be

used to determine reliability of hue determination. Low saturation indicates that the pixel is nearly gray, and as such it's color will be more variable than a 'strongly colored' pixel. François and Medioni use this approach with good results in [1], ignoring the hue difference when the saturation is below some threshold. However, thresholding discards information which could be useful in making a final decision, such as when two mostly unsaturated but nonetheless differently colored objects have the same intensity. The use of a saturation threshold also creates the requirement of either having a universally applicable threshold or resetting the threshold for different situation, which is undesirable.

I attempt to describe a method to adaptively determine a weighting function which can be used in combining the hue and intensity difference information as a continuous function of the reliability of a pixel's hue. Section 2 will discuss an iterative algorithm for weight modification, the results of which will be discussed in section 3. Conclusions will follow in section 4.

2. Procedure

The current system uses a log-scaled coloredness metric to classify pixels in a number of training images. It maintains the weighting function as two discreet log-scaled-coloredness-indexed lists of weights, one for hue and one for saturation. The system scans through pixels in the training data, comparing hue and intensity performance to manually segmented ground truth data. The error for each differencing operation is considered to be the difference between the desired output and the actual output. In each case, the weight for each differencing operation and for the particular coloredness of the pixel in question is adjusted by an amount proportional to the magnitude of the error. The development of the log-scaled coloredness metric and the training algorithm used are discussed further in the next subsections.

2.1. Coloredness Metric

Unfortunately, the commonly used definition of saturation is undefined for black pixels. Since even the smallest measurable amount of noise in a black pixel would cause wild changes in the hue angle, I require some measure of how colored a given pixel is which is well defined at and near black. One obvious choice for a metric of coloredness is the Euclidian distance from that pixel to the gray line which runs from white to black in RGB space (Figures 1 and 2).




Figure 1. Coloredness is the orthogonal distance from a pixel to the gray line.



Figure 2. Coloredness of the background image.

Unfortunately, the histograms of image coloredness in the test set, as well as a number of other sample images, tend to concentrate in the low values, with few strongly colored pixels. This makes it difficult to accurately determine a good weighting function in the upper range of the coloredness spectrum, as there are only a few training examples in this range. To remedy this histogram clustering, the coloredness was re-scaled using a natural log according to




Figure 3. Histogram of the background image coloredness.




Figure 4. Histogram of the log-scaled background image coloredness. The values are more evenly spread across the available range, increasing the resolution in the critical gray area.

$$scaled = \frac{(\ln(coloredness+0.9) - \ln(0.9))}{(\ln(1.9) - \ln(0.9))}, \text{ as seen in Figures 3 and 4.}$$

Although this is clearly not a perfect solution, it does stretch the histogram out somewhat, increasing discriminability in the gray range. This scaling was chosen because although it does not stretch the histogram very much, it does not cause much bunching at the high end of the coloredness spectrum, which was a problem with some of the other sample images and scaling equations.

2.2. Weight Adjustment

The weighting functions which determine how hue and intensity difference information are combined are initially set to a constant value for both hue and intensity and for all values of log-scaled coloredness. For this experiment, a set of hand segmented images representing a video conferencing scenario is used as the ground truth for training.



Figure 5. Hue difference. Random false positives around the shoulders correspond to gray areas where hue is unreliable.



Figure 7. Ground truth detection mask. The white silhouette should be detected as foreground, the black area as background.



Figure 6. Intensity Difference. There are virtually no false positives, and unfortunately the false negatives may be intractable.

3. Results

The results of the current system indicate its ability to detect the correlation between low log-scaled coloredness and unreliable hue differencing information (Figure 9), and to suppress hue detection information where the coloredness indicates that it will be unreliable. Figure 8 shows the result of combining the hue and intensity differences of Figures 5 and 6 using the weighting functions shown in Figure 9. Unfortunately, the homogeneity of the data set introduces noise into the weighting functions, resulting from repeated occurrence of difficult to detect false positives. Because the training data are similar, these difficult pixels don't average out and instead lead to correlated dips in the hue and intensity weights. Although the weighting does manage to suppress much of the spurious detection due to unreliable hue determination, it is not much better than the raw intensity difference at this point.

For each pixel in a training set, the weighted hue (Figure 5) and intensity (Figure 6) differences are calculated using weights determined by the pixel's coloredness, and then compared to the desired result (Figure 7). If the desired result is a detection (represented by a 1) and the calculated value was less than 1, the weighting function at that coloredness is increased by some amount proportional to the magnitude of the error. Similarly, if the desired result is background (denoted by a 0 in the ground truth segmentation) and the detection is greater than 0, the weighting function is decreased. All adjustments are scaled by a learning rate which reduces the possibility of harmful effects by a few rogue points, and by a histogram weight term which averages the results of all pixels in a particular log-scaled-coloredness-histogram bin.




Figure 8. Weighted Combination of hue and intensity differences.




Figure 9. Hue and intensity weighting curves as a function of log-scaled coloredness. Divergence near the origin indicates relative hue suppression. Correlation over the rest of the graph indicates that hue and intensity perform similarly for more strongly colored pixels.

4. Conclusions

Although the system has shown perceptible improvements in some cases, it has yet to find a weighting which consistently outperforms both hue and intensity differencing in a significant way. The infrastructure is nearly completed, and should perform better with a few more adjustments to the training algorithm and the use of a more tractable training set. Although clearly the goal is for the system to work on a wide variety of data sets, the next step of adaptively determining a simple parametric relationship between log-scaled coloredness (or some other color strength metric) and hue reliability should be simpler on a properly chosen training set which includes more examples of the problems to be overcome, such as shadows, and fewer intractable examples, such as identical background and foreground pixels.

References

- [1] A. R. Franois and G. G. Medioni. Adaptive color background modeling for real-time segmentation of video streams. In *Proceedings of the International on Imaging Science, Systems, and Technology*, pages 227–232, June 1999.
- [2] T. Horprasert, D. Harwood, and L. Davis. A statistical approach for real-time robust background, 1999.