

COMPUTER VISION SYSTEM FOR IDENTIFYING ROAD
SIGNS USING TRIANGULATION AND BUNDLE
ADJUSTMENT

by

ANUPAMA KRISHNAN

B.E., Anna University, 2007

A THESIS

submitted in partial fulfillment of the
requirements for the degree

MASTER OF SCIENCE

Department of Electrical and Computer Engineering
College of Engineering

KANSAS STATE UNIVERSITY

Manhattan, Kansas

2009

Approved by:

Major Professor
Dr. Chris Lewis

Copyright

Copyright (c) 2009 Anupama Krishnan Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

Abstract

This thesis describes the development of an automated computer vision system that identifies and inventories road signs from imagery acquired from the Kansas Department of Transportation's road profiling system that takes images every 26.4 feet on highways through out the state. Statistical models characterizing the typical size, color, and physical location of signs are used to help identify signs from the imagery. First, two phases of a computationally efficient K-Means clustering algorithm are applied to the images to achieve over-segmentation. The novel second phase ensures over-segmentation without excessive computation. Extremely large and very small segments are rejected. The remaining segments are then classified based on color. Finally, the frame to frame trajectories of sign colored segments are analyzed using triangulation and Bundle adjustment to determine their physical location relative to the videolog system. Objects having the appropriate color, and physical placement are entered into a sign database. To develop the statistical models used for classification, a representative set of images was segmented and manually labeled determining the joint probabilistic models characterizing the color and location typical to that of road signs. Receiver Operating Characteristic curves were generated and analyzed to adjust the thresholds for the class identification. This system was tested and its performance characteristics are presented.

Table of Contents

Table of Contents	iv
List of Figures	vi
List of Abbreviations	viii
1 Introduction	1
2 Segmentation	3
2.1 K-Means Segmentation	3
2.2 Two Pass Segmentation	4
2.3 Difference Image Segmentation	5
3 Classification	6
3.1 Elimination of Large and Small Segments	6
3.2 Color based Classification	7
4 Shape Classifier	10
5 Location Classifier	14
6 Camera Calibration	19
7 Triangulation	21
8 Bundle Adjustment	26
9 Optimal Threshold Value Selection	31
10 Receiver Operating Characteristic (ROC) Curves	39
10.1 ROC curve	39
10.2 Practical generation of ROC curves.	39
11 Results	42
12 Future Scope	47
13 Conclusion	48
Bibliography	51

List of Figures

2.1	K-Means Proper Segmentation	4
2.2	Under-segmentation	5
2.3	Difference Image Segmentation	5
3.1	Size Thresholding	6
3.2	Perfect Classification	8
3.3	Color Classifier	9
3.4	False positive	9
4.1	Fourier Descriptors	11
4.2	Shape Classification	12
5.1	4 Consecutive images	15
5.2	Trajectory	16
5.3	Consecutive images	17
5.4	Location Classification	17
6.1	Lens Distortion	20
7.1	Triangulation	23
7.2	High Triangulation Error	24
7.3	Curved motion	25
8.1	Test images	27
8.2	Objects in Scene	27
8.3	SIFT	28
8.4	Top View of Objects	28
8.5	Front View of Objects	29
8.6	Road Sign Size Estimation	30
9.1	PDFs of Hypotheses	32
9.2	Likelihood function	33
9.3	ROC for Not-White/White	35
9.4	Image with white road sign	36
9.5	$\alpha = 0.1$	37
9.6	$\alpha = 0.03$	38
10.1	ROC - Green color classification	40
11.1	Sample results set 1	43

11.2 Sample results set 2	44
11.3 Sample results set 3	45

List of Abbreviations

BA - Bundle Adjustment

pose - Position and orientation

SIFT - Scale Invariant Feature Transform

KDoT - Kansas Department of Transportation

ROC - Receiver Operating Characteristics

PDF - Probability Density Function

TPR - True Positive Rate

FPR - False Positive Rate

TP - True Positive

FP - False Positive

TN - True Negative

FN - False Negative

Acknowledgments

I would like to thank Dr. Dale Schinstock and Michal Jama for helping me understand the basic concepts involved in Bundle Adjustment. I would also like to thank Dr. Bala Natarajan for information on development of hypothesis tests which helped theoretically verify the research work.

Chapter 1

Introduction

This thesis describes development of an automated machine vision based road sign inventorying system that catalogs road signs from Kansas Department of Transportation's, (KDoT), imagery. KDoT's videolog system takes true color 1200x1600 pixel pictures every 26.4 feet along highways through out the state of Kansas. Many vision researchers have worked on building intelligent transportation systems addressing similar problems. Most utilize some means of segmentation process followed by a classification process that identifies the segments of a desired class.¹⁻⁵ Our system uses a novel difference image technique to efficiently achieve over-segmentation. It is also unique in that the 3D physical location of objects is ascertained and used to aid in classification. Being able to determine the 3-D location of imaged objects may permit database users to verify that road sign regulations are met as well as provide a myriad of other data mining opportunities. Our design is constrained by the massive number of images acquired yearly, and by the lack of control over lighting, viewing geometry and background. Algorithms are restricted to those that are computationally quick and memory efficient.

Processing begins by over-segmenting each image using two phases of a low cluster count K-Means algorithm. The first phase operates on the original image and the second on a difference image. Next overly large and small segments are rejected eliminating the vast majority from further consideration. Next, a statistical measure is used to classify segments as potential signs based on their color. Next, feature matching is performed to identify

sign colored segments that recur in consecutive frames. Recurring features are triangulated to estimate their 3D location relative to the vehicle's path. In the first triangulation pass, straight line motion between frames is assumed. The residual triangulation error is evaluated to determine if trajectory adjustment is warranted. When the error is large, usually on turns, Bundle Adjustment(BA) is used to refine motion estimates and facilitate better localization. Finally, another statistical classification process is performed using the 3D location. This technique takes advantage of the fact that signs are usually placed adjacent to the road at prescribed heights and distances from the road. A Neyman-Pearson hypothesis test is used to justify the use of Receiver Operating Characteristic curves for adjusting classifier thresholds. This thesis details the segmentation and classification process and presents results from an example set of highway images to demonstrate the efficiency and performance of the system.

Chapter 2

Segmentation

2.1 K-Means Segmentation

The goal of segmentation is to divide an image into clusters of like pixels. Ideally each cluster represents a single physical object. Like many image processing systems, we desire a segmentation that is slightly over segmented such that segments do not span multiple physical objects, but sometimes physical objects form multiple segments. K-Means is a computationally quick iterative algorithm that finds K natural clusters. It has seen widespread use for image segmentation for example, in medical applications such as in brain tumor detection⁶ and medical imaging⁷. The K-means algorithm iteratively finds the centers for k natural clusters in a set of data. The procedure for K-Means consists of the following five steps⁸:

Step 1: Initialize cluster centers $\mu_1, \mu_2, \dots, \mu_k$ and number of required clusters, k

Step 2: Classify each sample based on the nearest μ_i

Step 3: Recompute μ_i as the mean of the samples in each of the new classifications obtained from step 2

Step 4: Repeat steps 2 and 3 until there is no change in μ_i

Step 5: Return $\mu_1, \mu_2, \dots, \mu_k$

The number of clusters, k , is critical. Larger k results in more segments and increased computational complexity, while smaller values results in under-segmentation, but are computationally quick. For the images obtained from KDoT, sixteen clusters provided an appropriate trade off between adequate segmentation and computational complexity.

2.2 Two Pass Segmentation

Fig. 2.1, depicts the results with $k = 16$ with all pixels in each segment's color shifted to the closest cluster center's color. Even with only 16 colors, the image represents the original well. Note that the yellow road sign is only slightly shifted in color since k-means selected its shade of yellow as one of the 16 cluster centers.



Figure 2.1: *This figure shows the original image and the color segmented image obtained using K-Means Clustering on an image with a yellow road sign. In the color segmented image the pixels belonging to the yellow portion of the road sign are grouped into a segment.*

However, under-segmentation also occurs as shown in Fig. 2.2 where k-means assigns the small road sign to the cluster color associated with similarly colored background objects resulting in the merging of the two physical objects into one segment. This occurs often with green or white signs as vegetation or white sky are common backdrops. This deficiency is alleviated using a difference image technique.



Figure 2.2: Here, the green sign's color is not one of the 16 color centers. Its pixels are shifted to the closest cluster center resulting in under segmentation.

2.3 Difference Image Segmentation

Usually in cases where under-segmentation occurs with K-Means, smaller objects undergo large shifts since none of the cluster centers are that close to the object's original color. This phenomenon led to the use of a difference image which is formed as the euclidean distance in color space between the original image and the K-Means color segmented image for each pixel. The difference image accentuates objects receiving large color shifts. Proper segmentation of these problematic objects is achieved by reapplying K-Means to the difference image. Proper segmentation of the green sign from Fig. 2.2 shown in Fig. 2.3.




Figure 2.3: Significant color shifts are accentuated in this gray scale difference image, (left). K-Means applied to the color difference image properly segments the green sign.

Chapter 3

Classification

3.1 Elimination of Large and Small Segments

A large number of segments are obtained from the segmentation phase. Most of them are very small, a few pixels in size. Others are very large, such as segments of the sky. Neither need further consideration and are eliminated by applying size thresholds. For the current process, all segments that are either smaller than 50 pixels or larger than 20000 pixels are removed. Fig. 3.1 shows the segments remaining from Fig. 2.2 after thresholding.




Figure 3.1: *Black and white depiction of remaining segments after size threshold is applied to the color segmented image shown in Fig. 2.2 and difference image shown in Fig. 2.3.*

3.2 Color based Classification

Road signs are more uniformly colored than typical objects. They have a unique correlation between their red, green and blue components even with differences in lighting. The joint color distributions for different sign types were determined by manually extracting numerous road signs of every type from a set of images. To help automatically classify segments its proximity to each of these joint color distributions is determined using the Mahalanobis distance measure.⁹⁻¹¹ The Mahalanobis distance is given by Eq. (3.1).

$$r = \sqrt{(\vec{x} - \vec{\mu})^T \Sigma^{-1} (\vec{x} - \vec{\mu})} \quad (3.1)$$


where $\vec{\mu}$ is the mean vector and Σ is the covariance matrix of the prior joint color distribution and \vec{x} is an input point in multidimensional space⁸. For this research, \vec{x} contains the average values of the Red, Green and Blue components of all the pixels belonging to the segment being analyzed. An optimum distance threshold is practically determined from the knee of the Receiver Operating Characteristic (ROC) curves. These curves are generated using a test set of 25 images containing signs of each color. The knee of each of these curves corresponds to the threshold value that gives the best trade off between false positives rate and true positive rate. A Neyman-Pearson Hypothesis test provides the theoretical justification for setting thresholds in this manner.

Few natural objects have a brilliant color similar to yellow road signs and hence the color classifier is nearly ideal for identifying segments of this type as shown in Fig. 3.2. For other colors such a green and white signs, this is unfortunately not the case.

For example, Fig. 3.3 displays some reasonably sized segments having similar color to green road signs, but which are not signs. Another example of this case is illustrated in Fig. 3.4, where road stripes were falsely classified as road signs due to their similarity in color distribution. Therefore, additional classification steps are necessary to accurately identify most sign types.



(a)



(b)



(c)

Figure 3.2: *The color classifier is able to identify yellow signs with very high precision.*



Figure 3.3: Results of color based classification on segments obtained in 3.1. The highlighted regions depict the segments with a color distribution similar to that of road signs.



Figure 3.4: A case of false positives for yellow signs.


Chapter 4

Shape Classifier


Fourier descriptors are a common tool used for pattern recognition and shape discrimination.¹² They are determined using the periphery of the object. The advantage of using the Fourier descriptors is their invariance to the starting point of the boundary, scale and rotation.^{13,14} Fourier Descriptors begin by representing the boundary of a region as a sequence of complex numbers. By taking the Fourier transform of this sequence and retaining only the lower frequencies, an approximation of the shape can be obtained. For recognition purposes the ratio of the magnitude of the upper frequencies with the first frequency, provides a set of scale and rotation invariant features. Since, signs are generally simple polygons with few sides; only a few terms are required to represent their shape. For this work, the number of fourier descriptors chosen was 6. An example of the shape approximation is shown in Fig. 4.1. The blue boundary shows the best fitting circle, red shows the best fitting ellipse and magenta shows the shape approximation with six Fourier descriptors for the two segments.

The Mahalanobis distance between the computed Fourier descriptors for a segment and the distribution of the descriptors of the desirable shapes is calculated in a similar manner as in color classification. Again, a threshold is applied to discard segments with the wrong shape. The shape thresholds are also determined using the ROC analysis. The output after elimination using this shape classifier is shown in Fig. 4.1 and Fig. 4.2.

As fourier descriptors are strictly dependent on the boundary, they fail in instances where there are slight variations such as the case shown in Fig. 4.2. When tested on a




(a)




(b)

Figure 4.1: This figures shows the results of classification based on shape of segments remaining from color classification in the example shown in Fig. 3.3. Fourier descriptors describe shape from the boundary of the segment. The best fitting circle (Blue), ellipse (Red) and a 6th order shape (Magenta) are shown for the two segments. The shape classifier applied to both sets of correctly colored segments, eliminates segments with improper shape but identifies the sign with the desired shape.



(a) Image with green road signs, left, and resulting segments after color classification, right.



(b) Magnified image from (a) showing that one sign has an irregular boundary, (left). The shape classifier was able to identify the sign with the rectangular boundary but however misclassifies the sign with the irregular boundary as not sign.

Figure 4.2: *A case of false negative while using fourier descriptors for classification.*

set of 50 images, the shape classifier was observed to be 92% efficient at identifying road signs correctly but misclassified 8% of the road signs as not signs. This shortcoming was overcome by replacing the shape classifier with one based on the feature's frame to frame trajectory.

Chapter 5

Location Classifier


Images are captured sequentially every 26.4 feet. Therefore, road signs are likely to appear in a sequence of several images and follow a characteristic trajectory. Matching features from segments appearing in consecutive images and comparing their trajectory through the field of view of the camera to that typical of road signs eliminates segments that are not road signs. In our system, the SIFT algorithm,¹⁵ tracked features from frame to frame.

The 2-D trajectory of the features identified by SIFT was compared to a probabilistic model determined from a test set of sequential images containing road signs. Similar to color based classification where the feature was a 3 dimensional vector, components of which were the red, green and blue values of the pixels, the feature space for the location based classification are four dimensional vectors defining the motion of typical road signs. The components of each vector are the u and v pixel coordinates of the road sign and the change in location in u and v of that matched feature in terms of pixels in the consecutive image. An example illustrating the trajectory of a road sign is shown in Fig. 5.1 and Fig. 5.2.


This process removed problematic segments including billboards containing similar color distribution. Fig. 5.3 shows the resulting segments from the size and color based classification on an example pair of consecutive images. The outcome of location based classification on these segments is shown in Fig. 5.4. Typically, all the false positives from the color classifier are removed by the trajectory classifier. Frame to frame feature matching is an



Figure 5.1: *Four consecutive images obtained from KDoT's road profiler*



(a) Trajectory



(b) Magnified trajectory

Figure 5.2: *Positions of road sign in four images*



Figure 5.3: *Two subsequent images obtained from the video log. The highlighted regions indicate the segments that remain after size and color based classification.*



Figure 5.4: *Results of location based classification. All the segments that do not follow the trajectory of a typical road sign are removed.*

unstructured and computationally efficient form of triangulation.

This method worked well on straight roads as shown in Fig. 5.3 and Fig. 5.4, but failed on curves where trajectory of objects is altered through the field of view. The system now uses multiple images to triangulate objects and ascertain their three dimensional location relative to the vehicle. These 3D positions help distinguish potential road signs from objects not correctly positioned. An initial triangulation is carried out assuming straight line motion but the residual error is monitored to detect non-linear vehicle motion. In those relatively infrequent cases where the vehicle turns substantially, the vehicle's true motion is estimated using a process called Bundle Adjustment (BA), and the 3D locations are re-estimated for

the classifier.


To implement a 3D location classification, features must be tracked frame to frame. As the vehicle moves, images of objects change in size. The Scale Invariant Feature Transform (SIFT) algorithm¹⁵, is used by our system to track objects frame to frame because it is robust to changes in scale, rotation and lighting. SIFT is state-of the-art, and therefore commonly applied to video tracking applications.^{16,17} An example of feature matching using SIFT is shown in Fig. 8.3. After identifying objects that appear in successive images using SIFT, it is important to remove lens distortion from the images to ensure a more accurate triangulation. The next chapter describes the procedure that was adopted to remove the distortion from the regions of interest identified by SIFT.

Chapter 6


Camera Calibration

Real lenses distort images and corrupt the triangulation process. For our purposes, it is not necessary to undistort the entire image. Rather, to minimize computations, lens distortion effects are taken into account only for the features tracked frame to frame. The imaging model for a camera including lens distortion is determined through camera calibration.¹⁸ Calibration parameters include the focal lengths, f_x and f_y , the principal point, (c_x, c_y) , and the skewness.¹⁹ To estimate these parameters, this effort used a Calibration toolbox for Matlab to process multiple images of a large checker board pattern printed on a 4x8 sheet of structural foam. The calibration process is explained in the work by Zhang¹⁸ and Heikkila²⁰ and the tool itself is available on line.²¹ To permit accurate triangulation, these calibration parameters are used to shift the location of features to where they would appear in the focal plane of an ideal pin hole camera. Photographs of the calibration target, with and without distortion, are shown in Fig. 6.1.

The next chapter describes the triangulation process adopted to address the research problem.



(a) With lens Distortion



(b) Without lens Distortion

Figure 6.1: Checkerboard pattern used to calibrate for lens distortion, top. Unwarped image shown on the bottom.

Chapter 7

Triangulation

Triangulation is a well known photogrammetric process for determining the 3-D location of features that appear in two or more images taken from different vantage points. Several methods exist for triangulation²². The optimal triangulation method involves the evaluation of the roots of a sixth degree polynomial to compute the best estimate²². However, this is computationally expensive and hence a very computationally simple method was employed to obtain an estimate for the 3-D location of features in a pair of successive images. We use an analytic solution where the location estimate is the midpoint of the line segment passing between the points of closest approach between two rays, one from each camera. The pin hole camera model depicted in Fig. 7.1 defines each ray, \vec{q}_1 and \vec{q}_2 , emanating from the focal points of each camera, passing through the image plane at the pixel coordinates u, v where the feature was observed in each image. Ideally these two rays intersect at the feature. Noise in camera position, orientation or in lens distortion results in these rays being skew. The shortest line segment between the two rays is mutually perpendicular to both and its direction is defined by their cross product as follows,

$$\vec{q}_3 = \vec{q}_1 \times \vec{q}_2 \quad (7.1)$$

A vector loop equation involving these rays defined as follows

$$\vec{C}_1 + \gamma_1 \vec{q}_1 + \gamma_3 \vec{q}_3 = \vec{C}_2 + \gamma_2 \vec{q}_2 \quad (7.2)$$

This may be rewritten as:

$$\vec{C}_1 - \vec{C}_2 = [-\vec{q}_1 \ \vec{q}_2 \ -\vec{q}_3] \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \end{bmatrix} \quad (7.3)$$

and then solved to determine the unknowns γ_1 , γ_2 and γ_3 . The midpoint of the line segment

$$\vec{f}_{xyz} = \vec{C}_1 + \gamma_1 \vec{q}_1 + \frac{\gamma_3}{2} \vec{q}_3 \quad (7.4)$$

is our analytic triangulation estimate. Back projecting this midpoint into both images and comparing to the observed image provides the measure of residual triangulation error as follows:

$$e = (u_1 - \hat{u}_1)^2 + (v_1 - \hat{v}_1)^2 + (u_2 - \hat{u}_2)^2 + (v_2 - \hat{v}_2)^2 \quad (7.5)$$

where \hat{u} and \hat{v} are back projected image estimates given by:

$$\begin{bmatrix} \hat{u}_i \\ \hat{v}_i \end{bmatrix} = \frac{f}{z} \begin{bmatrix} x \\ y \end{bmatrix} \quad (7.6)$$

where x, y , and z are the coordinates of f_{xyz} , expressed the i^{th} camera's frame.

The initial triangulation is performed assuming straight line motion without any rotation because images are usually taken on straight highways with little curvature. In this case the difference between the vectors \vec{C}_1 and \vec{C}_2 is accurately known. However, on turns, the straight line assumption is not met and triangulation error is high as illustrated in Fig. 7.2. For these cases, the motion estimate is refinement using BA.




Figure 7.1: A feature is imaged by two cameras located at \vec{C}_1 and \vec{C}_2 . The rays \vec{q}_1 and \vec{q}_2 from these locations are defined by the feature's u and v pixel coordinates and the associated focal lengths f_1 and f_2 . Ideally, these two rays intersect at the feature, but are more typically skew. Triangulation is the process of estimating the feature's location \vec{f}_{xyz} .

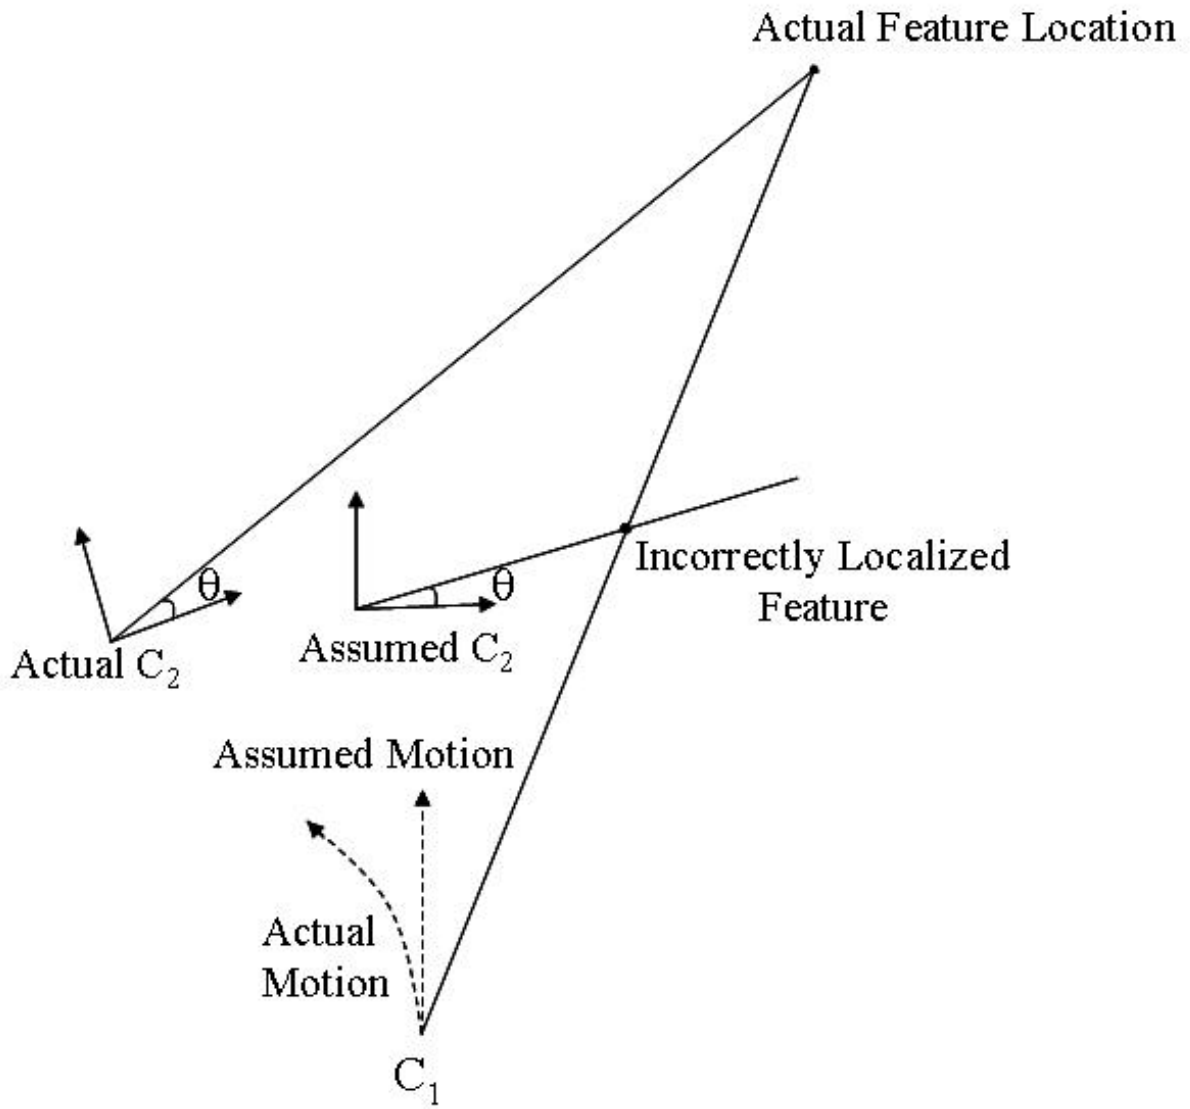


Figure 7.2: *Triangulation error is large during turns.*



(a)



(b)

Figure 7.3: *Usually images are taken with straight line motion between frames, and the vectors \vec{C}_1 and \vec{C}_2 are accurately known. However, on turns, the straight line assumption is not met and triangulation error is high.*

Chapter 8

Bundle Adjustment

Bundle Adjustment, (BA), modifies the camera's pose estimate including position, \vec{C}_i , and the orientation, (roll (R_r), pitch (R_p) and yaw (R_y)) along with the 3-D locations, $T_i = (x_i; y_i; z_i)$, of the features to reduce the triangulation error. In our implementation, the frame of reference is the first camera's frame. The six pose parameters of the second camera and three location parameters for each feature relative to this frame are optimized. The vector of parameters for m tie points is given by: $\vec{K} = [\vec{C}_2^t; R_r; R_p; R_y; T_1; \dots T_m]$. The cost function optimized by BA is the sum of triangulation errors over all features. This non-linear cost function is solved with the iterative Levenburg-Marquardt algorithm²³. This thesis only presents the basic idea behind Bundle Adjustment. Further elaboration on the concept of BA and comparison between different cost functions is explained by Triggs²⁴ and by Hartley and Zisserman²². Without constraints other than tie points scale is not determined. Constraining our solutions to match the known $\|\vec{C}_1 - \vec{C}_2\| = 26.4$ feet of motion between images fixes scale and for our system BA converges providing initial pose estimates are reasonably accurate. Fig. 8.1 to Fig. 8.6 depicts the results of using triangulation and BA on some objects that appear in a pair of successive images obtained from KDoT's videolog system. From the pair of images, the relative locations of various objects is determined. The physical location of objects is strongly tied to whether or not they are a sign.

Classification based on physical location is analogous to color based classification. To characterize the statistical distribution of sign locations, a database was created from 50



Figure 8.1: *Two consecutive images to be analyzed.*




Figure 8.2: *Scene with various objects identified for triangulation. BA will reduce triangulation error and allow scene reconstruction.*




Figure 8.3: For triangulation, features must be matched from frame to frame. SIFT is used to accomplish this.

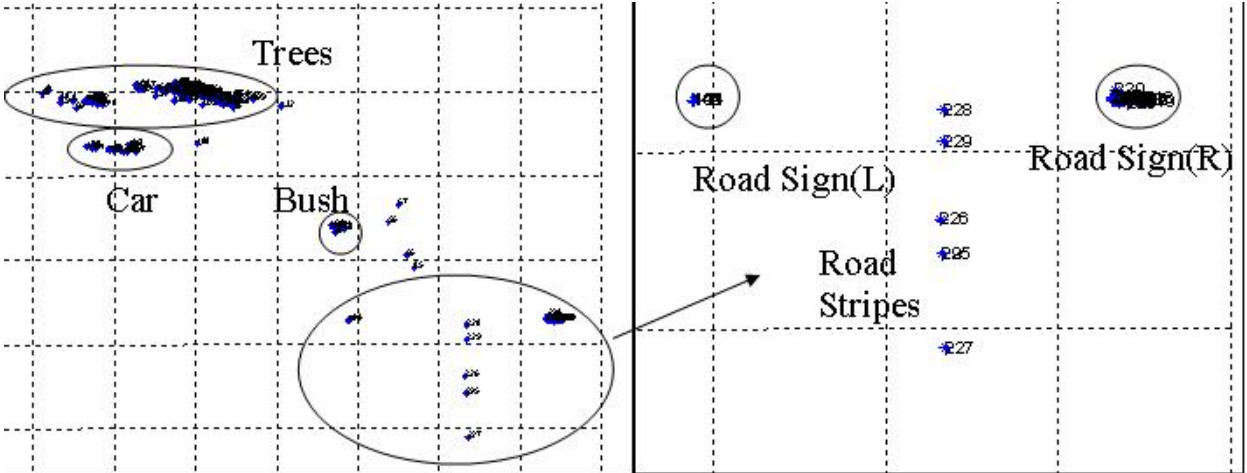


Figure 8.4: The top view of the 3-D scene containing the labeled objects obtained from BA. Note how BA accurately computes the 3-D location of the features extracted from the road stripes. They appear in between the yellow road sign on the left and the green and white road signs on the right.

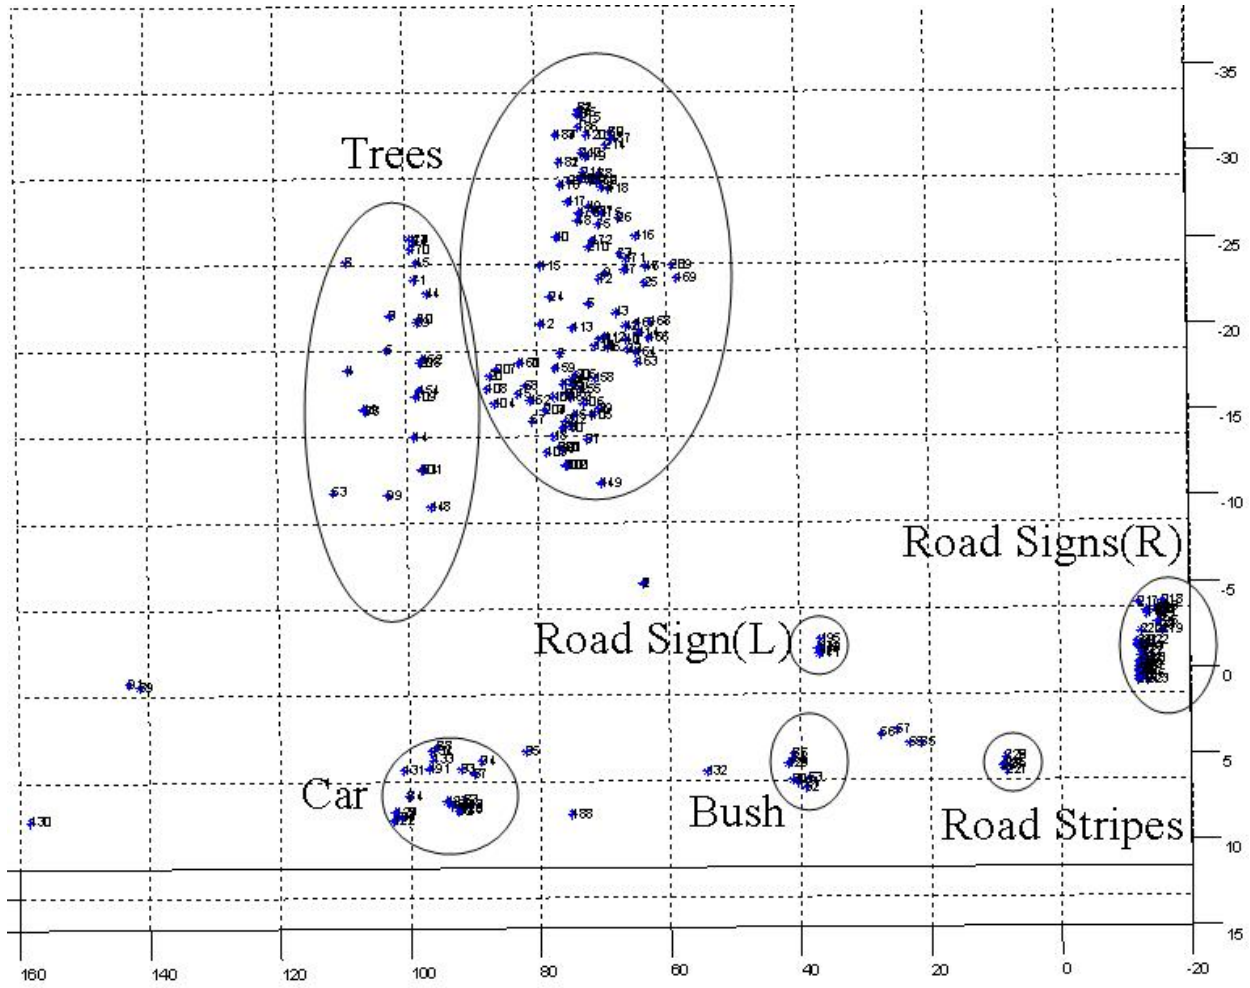


Figure 8.5: *Front View of 3-D plot of the locations of the matched features computed by BA.*

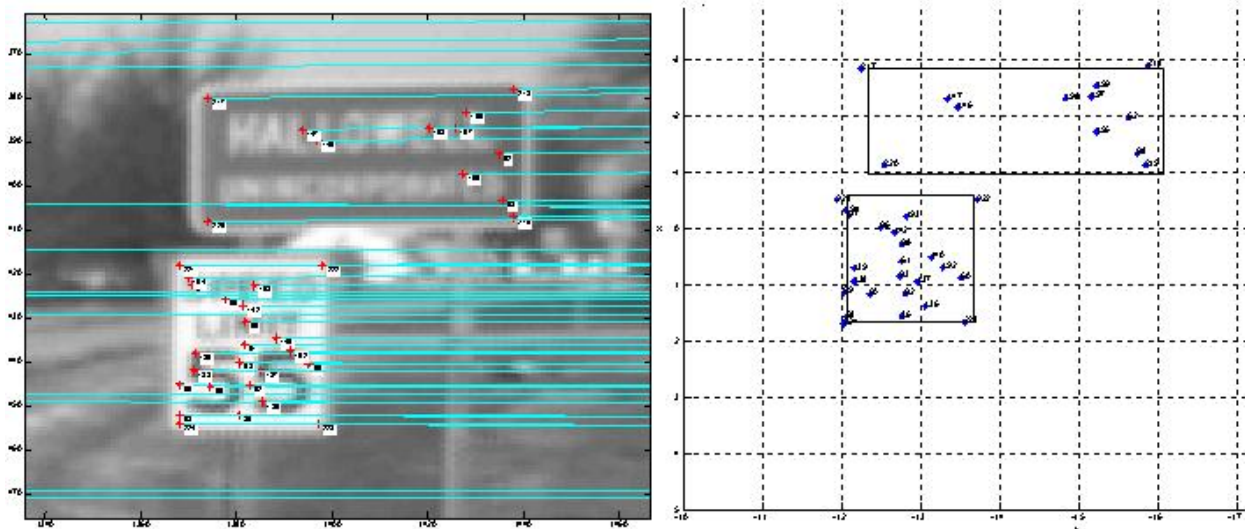


Figure 8.6: Features matched from the two road signs on the right are shown on the left and the zoomed view of figure 9 showing the the front view of the corresponding 3-D location of the features obtained using BA is shown on the right. Each division is 1 foot and it can be seen that the larger road sign is about 3.5 feet wide and has a height of about 2 feet. Hence an estimate of the size of the road sign is obtained.

images containing the 3-D location of features from road signs. Tie points were extracted using SIFT on just the regions of the image containing the road signs. Their 3D locations were determined using triangulation and BA. The mean (x, y) values of the positions of road signs was $(1.9783, 10.2391)$ in the first camera's frame, ten feet to the right and two feet above the camera. Note, since z is along the direction of travel, that it cannot be used. The Euclidean distance from this mean is now the Mahalanobis distance measure used for location based classification.

The Mahalanobis distance measures the proximity of a point to a distribution. This measure is used in both the color and location classifiers. Distance thresholds need to be determined for decisions to differentiate between objects being signs or not. The basis for selecting these critical threshold values is described in the next chapter.

Chapter 9

Optimal Threshold Value Selection

The selection of thresholds for classification can be developed as a hypothesis testing problem. The details on the various types of hypothesis tests and their comparisons are given by Poor²⁵. Here we apply the Neyman-Pearson test with a binary hypotheses for detecting white road signs. The problem can be defined as:

H_0 : Not White Road Sign

H_1 : White Road Sign.

The statistics of color are modeled using 50 images to visually construct a database of the red, blue and green values of all the pixels in objects that are and are not white road signs. Models of probability density functions are fit to the color histograms for both hypotheses. The PDFs are illustrated in Fig. 9.1. For red data, the PDF is modeled as:

$$p_0(y) = \frac{1}{\sqrt{2\pi * 1489}} \exp \left\{ \frac{-(y - 118.337)^2}{2 * 1489} \right\} \quad (9.1)$$

$$\forall \quad 20 \leq y \leq 252.6$$

$$p_1(y) = 4.211 * 10^{-9} \exp \left\{ \frac{y}{\sqrt{235}} \right\} \quad (9.2)$$

$$\forall \quad 136.2 \leq y \leq 253.8$$




Figure 9.1: *PDFs for the two hypotheses: $p_0(y)$ and $p_1(y)$*

The likelihood ratio test is defined by:

$$\text{choose} \begin{cases} H_1 & \text{if } L(y) \geq \eta \\ H_0 & \text{if } L(y) < \eta \end{cases} \quad (9.3)$$

where $L(y)$ is the likelihood ratio function given by:

$$L(y) = \frac{p_1(y)}{p_0(y)} \quad (9.4)$$

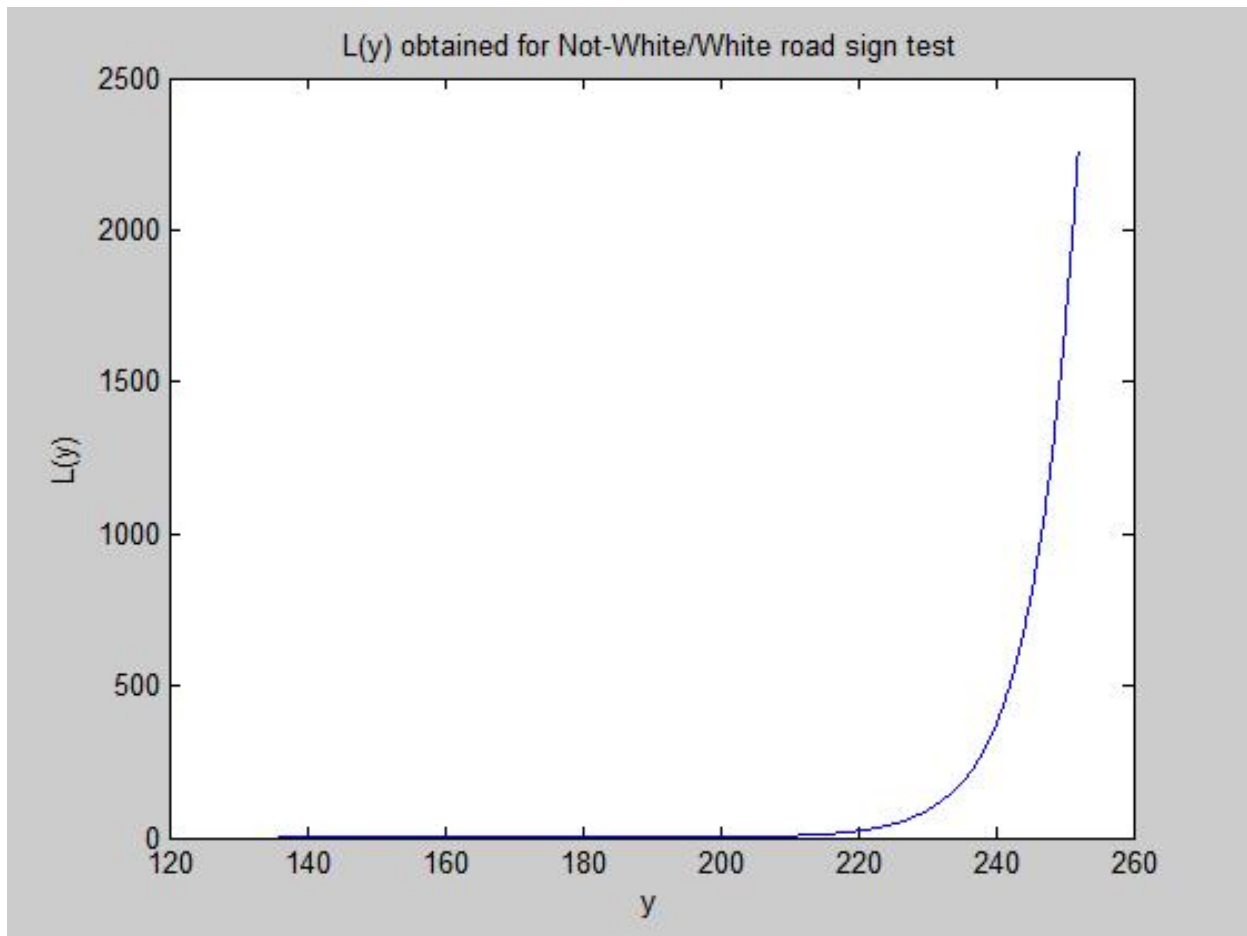


Figure 9.2: *Monotonic and continuous Likelihood function.*

Since in this case $L(y)$ is continuous and monotonic, as shown in Fig. 9.2, the relationship between $L(y)$ and η can be expressed as a relationship between y and some threshold y_η . Randomization is not necessary. Therefore, the decision rule for the Neyman-Pearson test

is:

$$\delta(y) = \begin{cases} 1 & \text{if } y \geq y_\eta \\ 0 & \text{if } y < y_\eta \end{cases} \quad (9.5)$$

The trade-off between probability of false alarm and probability of detection is attained by setting a bound on the probability of false alarm.

$$\max_{\delta} P_D(\delta) \text{ subject to } P_F(\delta) \leq \alpha \quad (9.6)$$

Where $P_D(\delta)$ is the probability of detection, $P_F(\delta)$ is the probability of false alarm and α is the bound mentioned earlier and it is also known as the level or significance level of the test²⁵. Probability of false alarm is computed as a function of y_η by finding the area under the curve $p_0(y)$ in the observation space of hypothesis H_1 . This expression is equated to α . The decision threshold y_η is then expressed in terms of the false alarm α . P_D is then calculated as the area under the PDF $p_1(y)$ in the observation space of H_1 . The expressions obtained after setting $P_F = \alpha$ for y_η and P_D are given as:

$$y_\eta = 118.337 - 54.57 \operatorname{erf}^{-1}(2\alpha - 0.9995) \quad (9.7)$$

$$P_D = 6.4555 * 10^{-8} \left(15495907.18 - \exp \left\{ \frac{y_\eta}{\sqrt{235}} \right\} \right) \quad (9.8)$$

The plot of P_D Vs P_F as shown in Fig. 9.3 is called the Receiver Operating Curve(ROC).

The effect of change in the value of alpha is illustrated in Fig. 9.4, Fig. 9.5 and Fig. 9.6. Increasing values of α lead to more false positives. Setting the threshold at the knee of the curve gives the best trade off between the probability of false alarm and probability of detection.

Hypothesis tests were not developed for all colors nor the location classifiers as some of the distributions were more complex. Instead, thresholds determined by generating ROC curves directly from the test data and setting the thresholds to the knee of the ROC curve. The next chapter describes how the ROC curves were estimated experimentally.

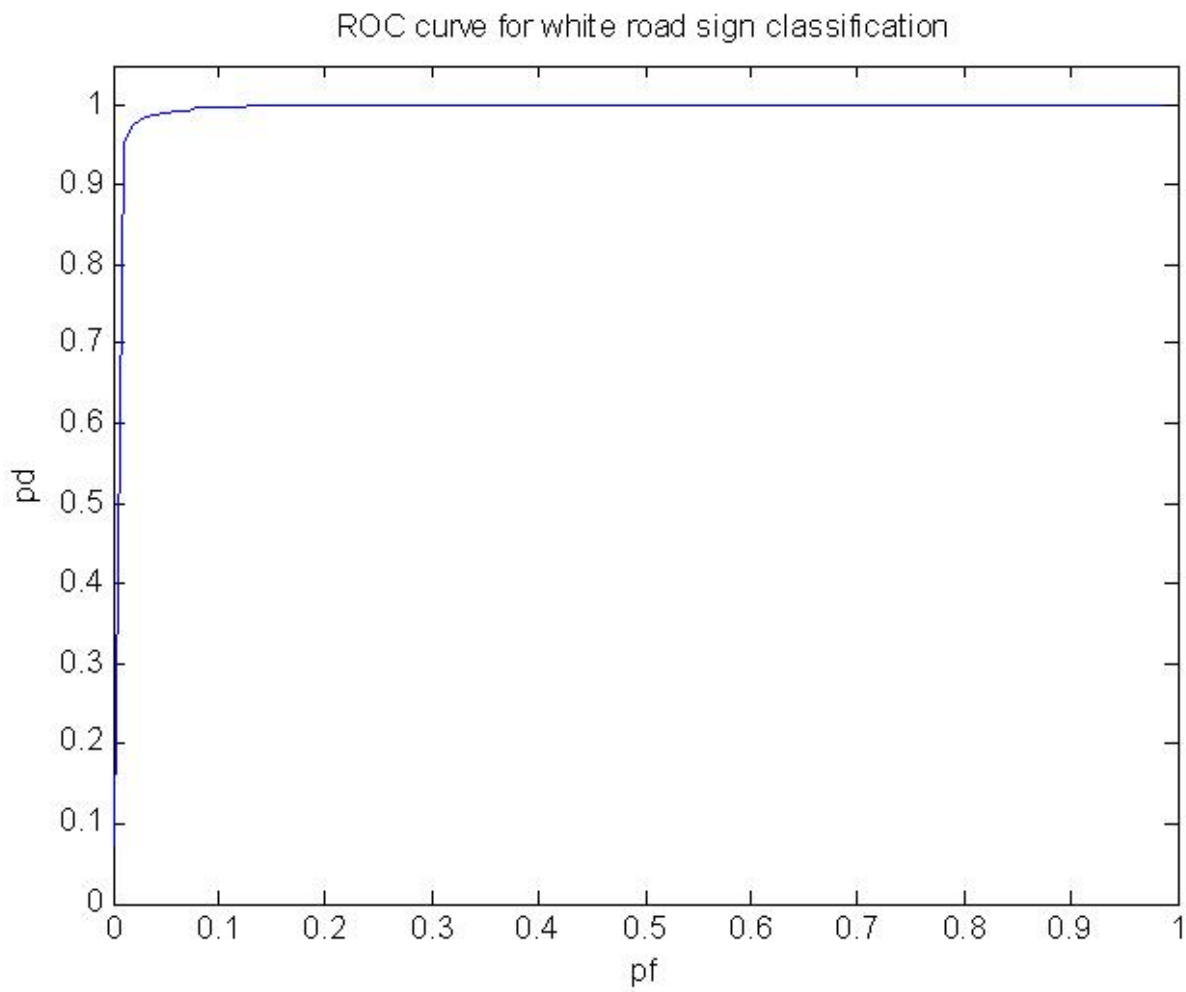


Figure 9.3: ROC obtained for the test designed for Not-White/White road signs based on color



Figure 9.4: *An image obtained from KDoT's videolog system.*



Figure 9.5: $\alpha = 0.1, y_\eta$ obtained is 172. More false positives. All the segments that have pixels whose Red, Green and Blue values are above y_η are classified as white road signs.




Figure 9.6: $\alpha = 0.03$. y_η obtained is 207. Lower value of alpha results in fewer false positives.

Chapter 10

Receiver Operating Characteristic (ROC) Curves

10.1 ROC curve

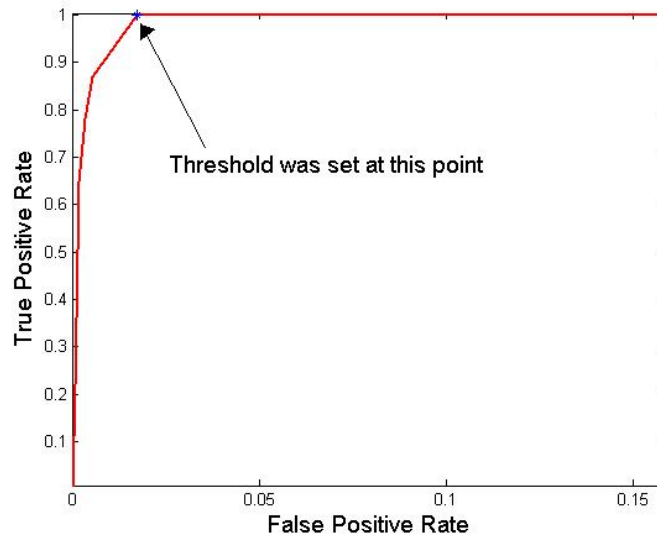
Classifiers produce four possible outcomes. If the classifier correctly identifies an object, it is a true positive (TP). If an object is incorrectly classified it is a true negative (TN). Objects of a class label not classified as such represent false negatives (FN). Likewise, objects not of the given type that are falsely identified as such are false positives (FP). The confusion matrix is shown in Fig. 10.1b. ROC curves display the relationship between True Positive Rate (TPR) and the False Positive Rate (FPR) as the threshold is adjusted. The closer the curve is to the top left corner the better the performance²⁶. The TPR and the FPR is calculated using:

$$TPR = \frac{TP}{TP + FN} \quad (10.1)$$

$$FPR = \frac{FP}{FP + TN} \quad (10.2)$$

10.2 Practical generation of ROC curves.

Using the database of manually labeled objects from 50 images, the ROC curves for both RGB color and XY distance thresholds and each sign color were constructed. An example



(a) ROC curve relating the classification threshold to the TPR and FPR for the color classifier

		Actual Value	
		P	N
Predicted Value	P	True Positive	False Positive
	N	False Negative	True Negative

(b) Confusion matrix

Figure 10.1: The arrow at the knee of the curves indicates an optimal threshold. A higher threshold minimally improves the TPR, but significantly increases the FPR. Note, since color performs well for classification the x-axis scale is magnified.

ROC curve obtained for green colored road sign classifier is shown in Fig. 10.1a. In the example, with the threshold set at the knee as indicated all the segments belonging to road signs were identified and 1.8% of the non-sign segments were misclassified as road signs.

Chapter 11

Results

To check the validity of the algorithm, it was used to identify all green road signs from images acquired from two counties which jointly comprised of about 20000 road side images. The images were first over-segmented using K-Means clustering on first the original image and then on the difference image computed by taking the euclidean distance in color space between the original image and the color segmented images. In the next stage, the segments classified based on size. This was followed by classification of segments based on their color distribution. The Mahalanobis distance was used as the metric to determine the closeness of a segment's joint color distribution to that of typical road signs and ROC curves were used to determine the value for the decision threshold. The segments were then classified based on their 3-D physical location as being signs and the threshold for the location classifier was also determined from ROC curves. The identified road signs were recorded in a database.

There were 89 green road signs situated in both counties together and all of them were detected. The only cases when the road signs were not identified was when the road sign has either color faded or have been dislocated from their position due to some accident. Hence the main intent behind designing the algorithm, which is to identify these problematic signs, is fulfilled. Bundle Adjustment used up only about 0.0001% of the total computation required for processing the images for the two counties. Fig. 11.1, Fig. 11.2 and Fig. 11.3 show examples of some classification of different road signs from images acquired by KDoT's videolog system. BA was not required to aid in identifying road signs on curved roads as



(a)



(b)



(c)

Figure 11.1: *Sample results: Some results of road sign detection from road side images.*



(a)



(b)

Figure 11.2: *Sample results (cont.): Notice how in (a), due to color fading, the brown road sign is not detected. A record of this sign, hence, will not go into the road sign database indicating that it has to be replaced. With BA to estimate motion between frames, the triangulation and thus sign identification remains effective during turns (b).*



(a)



(b)



(c)

Figure 11.3: *More examples of road sign detection under varied lighting conditions. Note that the location classifier for the road signs on the left of the road is different to that of road signs on the right of the road.*

shown in Fig. 11.3. It was only in those rare cases where objects with the color and size typical to that of road signs appeared on consecutive images with sharp turns that BA was triggered. An example of a case where Bundle Adjustment was used to detect a signs that were right on a turn (Fig. 7.3) is shown in Fig. 11.2(b). Note that in Fig. 11.2(a), the brown sign had color faded and hence was not identified and recorded by the system indicating the need for replacing this road sign.

Chapter 12

Future Scope

The current method can identify road signs, determine their type, 3-D location with respect to the camera and size. Future scope of the research is to determination of the distance between road signs of a particular type and investigate processes that take advantage of 3D information ascertained from triangulation such as measuring road easement profiles. Once the road signs are obtained, the size of the road signs can be estimated just as in the example case shown in Fig. 8.6. Future work will also include identification of the type of post on which the road sign is mounted with the possibilities being wood post, metal post or steel beam. This can be implemented with the prior knowledge of the location of the road sign and the color information.

Chapter 13

Conclusion

This paper describes the development of an automated system that identifies and inventories road signs from imagery obtained from KDoT's videolog system of highways through out the state. There are two main stages of processing: Image segmentation followed by classification of segments identified. Initially, over-segmentation of images was achieved using K-Means clustering on first, the original image and then again on a difference image. The difference image was computed as the euclidean distance in color space between the original image and the color segmented images. This was followed by three stages of classification based on size, color and 3-D location of road signs relative to the camera. The size and color classifiers were adequate for yellow road signs as there are not that many naturally occurring yellow objects that have a color distribution typical to that of a yellow road sign. However, additional classification steps were warranted for classifying green and white road signs due to the background clutter such as trees and the sky. Triangulation and Bundle Adjustment(BA) were used to estimate the 3-D location of signs and a segment that had the right size and color was classified as a road sign if it was located at the right physical location with respect to the camera. Mahalanobis distance was used to measure the proximity of an object to being a road sign based on color and location. BA was used to refine location estimated during turns. ROC curves were used to identify decision thresholds for the color and location classifiers. The system was tested and the corresponding results are documented and explained. Future scope of this research is also presented.

Bibliography

- [1] Y. L. Andrzej Ruta and X. Liu, *Real-time traffic sign recognition from video by class-specific discriminative features*, Elsevier Journal of Pattern Recognition, Received 1st August 2007; revised 22nd May 2009; accepted 26th May 2009, Available online 8th June 2009.
- [2] C. Bahlmann, Y. Zhu, and M. Pellkofer, *A System for Traffic Sign Detection, Tracking, and Recognition Using Color, Shape, and Motion Information*, IEEE Intelligence Vehicles Symposium (*IV2005*), Las Vegas, 2005, p. 255-260.
- [3] H. Fleyeh and M. Dougherty, *Road and Traffic sign detection and recognition*, 10th EWGT Meeting and 16th Mini-EURO, 2005.
- [4] L. C. Gomez and O. Fuentes, *Color-Based Road Sign Detection and Tracking*, Image Analysis and Recognition (*ICIAR*), CA, 2007.
- [5] M. Benallal and J. Meunier, *Real-Time Color Segmentation of road signs*, Canadian Conference on Electrical and Computer Engineering *IEEE CCECE*, p. 1823-1826, 2003.
- [6] C.-C. C. Ming-Ni Wu, Chia-Chen Lin, *Brain Tumor Detection Using Color based K-Means Clustering Segmentation*, Third International Conference on International Information Hiding and Multimedia Signal Processing, *IIH-MSP*, Vol. 2, p. 245-250, 2007.
- [7] H. P. Ng, S. H. Ong, K. W. C. Foong, P. S. Goh, and W. L. Nowinski, *Medical Image Segmentation Using K-Means Clustering and Improved Watershed Algorithm*, IEEE Southwest Symposium on Image Analysis and Interpretation, p. 61-65, 2006.

- [8] R. O. Duda, P. E. Hart, and D. G. Stock, *Pattern Classification*, volume 2nd Ed., Wiley - Interscience Publication, 2001.
- [9] V. D and K. C, *Gaussian Mixture Modeling by Exploiting the Mahalanobis Distance*, *IEEE Signal processing Transactions*, 56, Issue 7, p. 2797-2811, 2008.
- [10] V. A. Kremens, R. L. Ononye, and A. T. .C, *A Hybrid Contextual Approach to Wildland Fire Detection Using Multispectral Imagery*, *IEEE Geoscience and Remote Sensing Transactions*, Volume 43, Issue 9, p. 2115-2126, 2005,.
- [11] V. F. N. C. E. V. Moraes. J. C. T. B., Seixas. M. O., *A real time QRS complex classification method using Mahalanobis distance*, *Computers in Cardiology*, p. 201-204, 2002.
- [12] E. Persoon and K. Fu, *Shape Discrimination Using Fourier Descriptors*, *IEEE transaction on systems, man, and cybernetics*, Vol 7, p. 170-179, 1977.
- [13] L. F. Costa and R. M. Cesar, *Shape Analysis and Classification, Theory and Practice*, CRC Press Boca Raton, Florida,, 2001.
- [14] I. Kunttu, L. Lepisto, J. Rauhamaa, and A. Visa, *Multiscale Fourier descriptor for shape-based image retrieval*, *pattern recognition, icpr*, vol 2, p. 765-768, 2004.
- [15] D. G. Lowe, *Distinctive Image Features from Scale-Invariant Keypoints*, *International Journal of Computer Vision*, p. 91-110, 2004.
- [16] T. T. Y. A. Akira Suga, Keita Fukuda, *Object recognition and segmentation using SIFT and Graph Cuts*, *ICPR Pattern Recognition*, p. 1-4, 2008,.
- [17] Z. Z. Xuelong Hu, Yingcheng Tang, *Video object matching based on SIFT algorithm*, *Neural Networks and Signal Processing*, p. 412-415, 2008.
- [18] Z. Zhang, *Flexible camera calibration by viewing a plane from unknown orientations*, *Seventh IEEE International Conference on Computer Vision*, p. 666-673, 1999.

- [19] Y. Sagawa, R.; Yagi, *Accurate calibration of intrinsic camera parameters by observing parallel light pairs*, *IEEE ICRA*, IEEE International Conference Robotics and Automation, p. 1390-1397, 2008.
- [20] S. O. Heikkila, J., *four-step camera calibration procedure with implicit image correction*.
- [21] J.-Y. Bouguet, *Camera Calibration Toolbox for Matlab*, [http : //www.vision.caltech.edu/bouguetj/calibdoc/](http://www.vision.caltech.edu/bouguetj/calibdoc/), 2008.
- [22] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, volume 2nd Ed., Cambridge University Press, 2003.
- [23] C. Buckley, *Photomosaicing and automatic topography generation from stereo aerial photography*, Master's thesis, Department of Mechanical Engineering, Kansas State University, 2008.
- [24] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon, *Bundle adjustment - a modern synthesis*, springer-verlag, 2000.
- [25] H. V. Poor, *An Introduction to Signal Detection and Estimation*, volume 2nd Ed., Springer Texts in Electrical Engineering, 1994.
- [26] T. Fawcett, *An Introduction to ROC analysis*, Pattern Recognition Letters, special issue on ROC analysis, Vol. 27, p. 861-874, 2005.

Appendix A

GNU Free Documentation License

Version 1.2, November 2002

Copyright (C) 2000,2001,2002 Free Software Foundation, Inc.

51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

1. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited

to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

2. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero

Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words. A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the

beginning of the body of the text.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgments”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

3. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

4. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document’s license notice requires

Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

5. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the condi-

tions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- (a) Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- (b) List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- (c) State on the Title page the name of the publisher of the Modified Version, as the publisher.
- (d) Preserve all the copyright notices of the Document.
- (e) Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- (f) Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- (g) Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- (h) Include an unaltered copy of this License.
- (i) Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified

Version as given on the Title Page. If there is no section Entitled “History” in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

- (j) Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the “History” section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- (k) For any section Entitled “Acknowledgments” or “Dedications”, Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgments and/or dedications given therein.
- (l) Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- (m) Delete any section Entitled “Endorsements”. Such a section may not be included in the Modified Version.
- (n) Do not retitle any existing section to be Entitled “Endorsements” or to conflict in title with any Invariant Section.
- (o) Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version’s license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

6. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment

to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgments”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements.”

7. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

8. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate,

or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

9. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgments”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

10. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

11. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit

to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document

under the terms of the GNU Free Documentation License, Version 1.2

or any later version published by the Free Software Foundation;

with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.

A copy of the license is included in the section entitled “GNU

Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this: with the Invariant Sections being LIST THEIR TITLES, with the

Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.