SAND REPORT

SAND2004-1867
Unlimited Release
Printed October 2004

Statistical Pressure Snakes based on Color Images

Hanspeter Schaub, ORION International Technologies

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@adonis.osti.gov
Online ordering:	http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.fedworld.gov
Online ordering:	http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2004-1867
Unlimited Release
Printed October 2004

Statistical Pressure Snakes based on Color Images

Hanspeter Schaub
ORION International Technologies
Albuquerque, NM 87106
schaub@vt.edu

Abstract

The traditional mono-color statistical pressure snake was modified to function on a color image with target errors defined in HSV color space. Large variations in target lighting and shading are permitted if the target color is only specified in terms of hue. This method works well with custom targets where the target is surrounded by a color of a very different hue. A significant robustness increase is achieved in the computer vision capability to track a specific target in an unstructured, outdoor environment. By specifying the target color to contain hue, saturation and intensity values, it is possible to establish a reasonably robust method to track general image features of a single color. This method is convenient to allow the operator to select arbitrary targets, or sections of a target, which have a common color. Further, a modification to the standard pixel averaging routine is introduced which allows the target to be specified not only in terms of a single color, but also using a list of colors. These algorithms were tested and verified by using a web camera attached to a personal computer.

Acknowledgment

Thanks to Christopher Smith from the University of New Mexico for the valuable discussions regarding the target error function of the statistical pressure snakes.

Contents

Figures

Intentionally Left Blank

Statistical Pressure Snakes based on Color Images

Introduction

Statistical pressure snakes (or snakes for short) are a method to analyze an image. A closed curve is parameterized through a set of control points which form the snake. The goal of the snake is to wrap around a specific image target. In practice, such snakes can be used to track targets moving along the image, as well as identify target properties such as size and shape. In traditional black and white images, the target might be defined in terms of a specific gray level. The snake will then outline a continuous contour in the image encompassing pixels with the target gray level. While computing the parametric snake curve, the control points are algorithmically adjusted along the image to attempt to satisfy various snake conditions, such as smoothness criteria and an equal-point-spacing criteria. An image pressure force is used to make the snake outline the image target. Various techniques have been established to include image pressure. The current work is based on statistical pressure snakes developed by C. Smith,[1] which utilize a pressure force function developed by Ivins and Porrill.[2] This pressure force function will make the snake explore an image and outline the target. It is not required to setup all initial snake points outside of the target, as is the case with the older image pressure methods.

The snake algorithm by C. Smith was found to provide a fast and reasonably robust solution. However, it only operated on a gray scale image. This snake routine was modified to operate on a full color image instead. The goal of this change is to obtain a new snake algorithm that is more robust to lighting changes and will be able to track an object even when portions of this object are subjected to very different lighting conditions. The HSV color space target selection algorithm is illustrated in Figure 1 where a door's wood color is being tracked.

Gray-Scale Pressure Force Computation

The statistical pressure snake routine needs to compute an image pressure force. This force pushes the snake control point outward (relative to the snake center) if the current mean pixel value at the control point is over an image target pixel. If not, then the snake control point has an image pressure force applied which will cause this point to contract. The

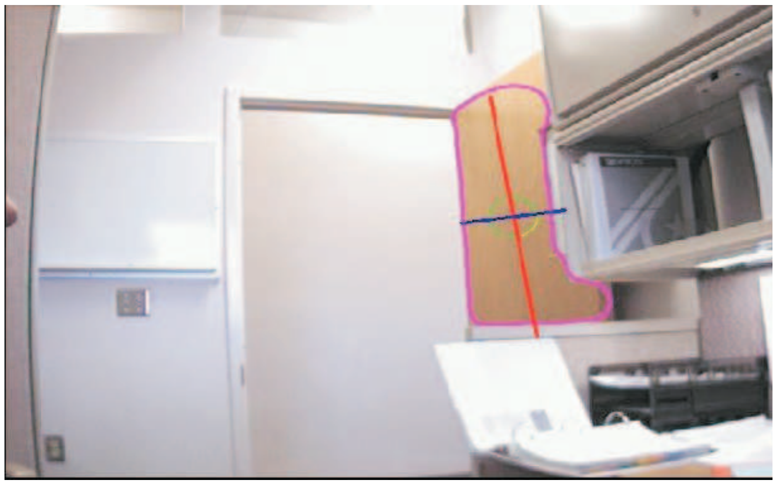

Figure 1. Statistical Pressure Snake Tracking the Wood Door

target pixels can be defined in many ways. In the original snake code, the target was simply defined as a gray scale value τ. The error signal ε of the current average pixel value p to the target value τ is then computed using

$$\varepsilon = \frac{|p - \tau|}{k\sigma} \tag{1}$$

where k is a user-chosen scaling factor and σ is the standard deviation of the target value τ in the image. The larger k is, the more forgiving the target value selection process is. If k is very small, then the target error signal will grow rapidly if the current mean pixel value is not equal to the target pixel value. The scalar target error signal ε is then processed in the snake code to produce the image pressure force for the current snake control point.

Color-Space Target Specification

Target Error Signal Computation

The pressure force calculation only requires a scalar measure indicating how well the current pixel values match those of the target. To make the snake routine work with the full color image, the target is specified in terms of all three color channels. Here the target error signal is defined as

$$\varepsilon = \sqrt{\left(\frac{p_1 - \tau_1}{k_1\sigma_1}\right)^2 + \left(\frac{p_2 - \tau_2}{k_2\sigma_2}\right)^2 + \left(\frac{p_3 - \tau_3}{k_3\sigma_3}\right)^2} \tag{2}$$

where p_i are local average pixel color channel values, τ_i are the target color channel values and σ_i are the target color channel standard deviations. The change to the pressure force subroutine code was minimal to incorporate a full 3 color channel image target specification.

Color Space Conversion

The key to making this color image error signal routine function well is the choice of color space to use. A standard Firewire web camera was used in developing this routine. The images captured are stored as a Red-Green-Blue (RGB) color image, where the three color channels contain the red, green and blue image components. This color space is commonly used by many cameras, since it emulates how the human eye perceives color. The incoming light is sensed by cones in the retina which measure the specific red, green and blue content of the light. With the web camera, an 8 byte number is used to store this information. A value of 0 means that no light was measured with the specific channel color, while a value of 255 means that a full value of that color was measured. A triad $(0,0,0)$ in RGB space provides a black color, $(255,255,255)$ is white, $(255,0,255)$ would provide purple, and so on.

While the RGB color space is routinely used when manipulating and printing images, it isn't ideal for the snake algorithm code. The goal is to have the target be identified under a range of lighting conditions. This means that if the target is specific type of red, for example, then we would like to register all shades of this red as the target color. It is possible to compute such a set of target colors in RGB space, but it is a very complicated function. Instead, it is convenient to map the image color channels from RGB space to Hue-Saturation-Value (HSV) color space. A similar color space is the Hue-Saturation-Intensity (HSI) color space. In the HSV color space the hue value determines the general color of the pixel (think of the color selection in a rainbow), the saturation value determines how rich or washed out the color is, and the value determines the brightness or darkness of the color.

The HSV color space is illustrated in Figure 2. Assuming the hue begins with red at 0, the hue values cycle through all possible colors until they return to blue at 360 degrees. For example, a pure green color would have a hue value of 120, while a pure blue would have

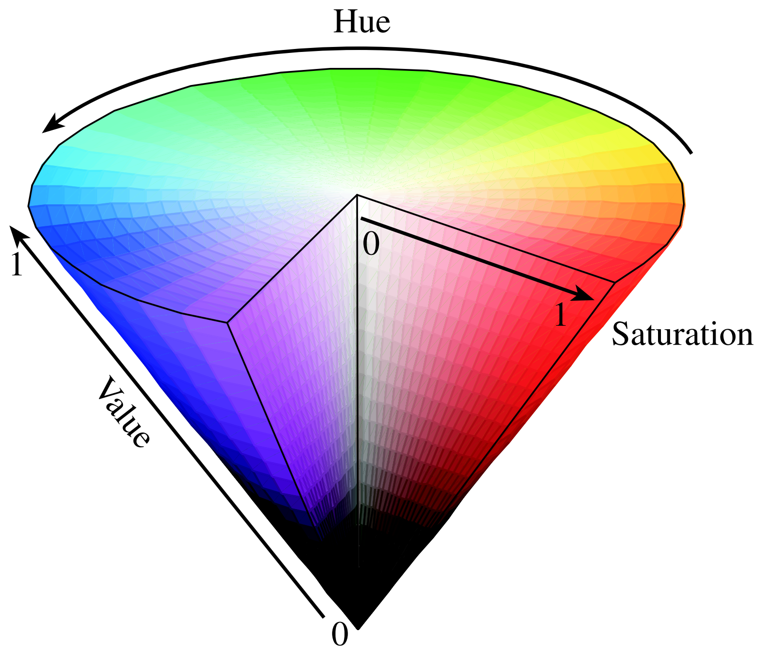

Figure 2. Illustration of the HSV Color Space

a hue of 240. If the saturation value S is 0, then the value V control the gray scale of the pixel. For gray scale images, the hue value is ill-defined.

Let us define our target color in the HSV color space. If we specify the target to only have a specific hue value (with associated standard deviation), then the target error function ε will return 0 if the current image pixels are of this hue type, regardless of how light or dark, saturated or wash-out the current pixel color is. All shades and variations of a color will be accepted. If the target is surrounded by a color with a clearly different hue (for example red box surrounded by a blue frame), then the snake will be able to easily track the red box, even tough the target might have drastically different lighting conditions across its surface. This idea is illustrated in Figure 3. Similarly, if portions of the box are more brightly lit, then these red zones with lower saturation value would also still be registered as the correct color. Such robustness to lighting changes between images, as well as across images, is not possible with classical gray-scale images. By moving to the HSV space, it is possible to make the snake routine a lot more robust to lighting changes as would be experienced in a typical unstructured outdoor environment (shadows being cast across object, clouds move in, etc.)

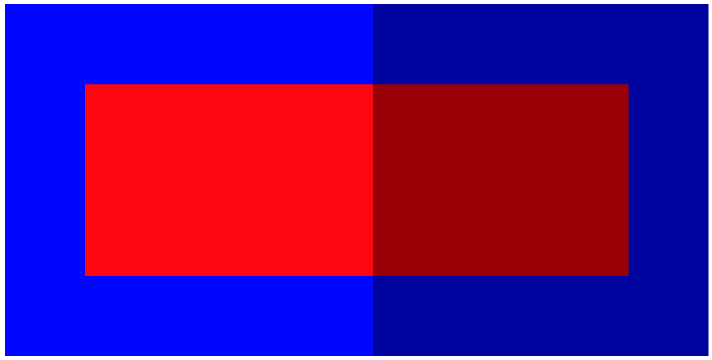

Figure 3. Illustration of shadow across Red Target Box

If the snake is to outline a general target, however, then only looking at the hue value to determine the error signal ε to the target color could be too forgiving. For example, consider that a blue suit-case is to be acquired. If only the hue value of a pixel is considered to generate ε, then a very white pixel, with only a slight hint of blue, would still be considered to be the target color. However, in practice this very light or dark color is often the color of another object in the scene. While the hue proximity is typically clipped with a gain of $k = 2$, it is found that for general target tracking the saturation values S and intensity values V should also be clipped with a gain of about $k = 6$. Thus, we are still weighting mostly on the hue parameter, but also penalizing for drastic departures from the target saturation and intensity values.

Using this HSV color space target selection scheme will have issues when the target color is near white or near black. For example, assume the target is nearly white, but has a slight shade of red to it. In this case it is possible that a different portion of the target might have a slightly blue shade (from a different light source, reflection off another target, etc.). The target selection scheme which mostly weighs the hue color component would have trouble here. If the target is mostly white (or black) a scheme which looks at the target saturation value would provide better results.

Multi-Color Snakes

Algorithm Changes

The previous algorithm assumed that the target was specified by a single HSV color. The target tracking error is computed by averaging the colors in a box of pixels about the current pixel of interest, and comparing this *averaged* color to the target color. However, if the target is specified by multiplied colors, this logic will fail at the border between two target colors. For example, consider the case where both blue (with a hue value of 240) and red (with a hue value of 0) are target colors. Next, assume the best case scenario where the box of pixels being averaged only contains these blue and red values. The numerical average will bet somewhere around 120, which corresponds to the color green. Thus, at this color border where both target colors are visible, the standard box averaging routine can result in a very large target error value.

The following algorithm modification will allow snakes to be defined in a multitude of colors, without causing the averaging problems at color borders. Instead of computing the required target error signal ε using the averaged pixel values p_i, we now compute an error signal ε_{ln} for each pixel in the 3x3 averaging box. Given the ith color channel image pixels m_{ijk} of pixel at (l,n), this intermediate error signal is computed using

$$\varepsilon_{ln} = \sqrt{\left(\frac{m_{1ln} - \tau_1}{k_1\sigma_1}\right)^2 + \left(\frac{m_{2ln} - \tau_2}{k_2\sigma_2}\right)^2 + \left(\frac{m_{3ln} - \tau_1}{k_3\sigma_3}\right)^2} \tag{3}$$

Applying this algorithm to each pixel for the 3x3 averaging box, we then compute an averaged target error signal ε for a snake point at (j,k) using

$$\varepsilon = \frac{1}{9}\sum_{l=j-1}^{j+1}\sum_{n=k-1}^{k+1}\varepsilon_{ln} \tag{4}$$

If the target is defined through multiple target HSV color sets, then the smallest error signal ε is used to drive the current snake point. For example, if the target color values are red and blue, and the local image color is red, then the error signal to the red target is very small, while the error signal to the blue target is very large. In this case, the error signal to the red target would be used to compute the image pressure term on the snake control point.

Numerical Results

This algorithm was implemented in the visual snake C-code. The results are shown in Figure 4. First, only a single target color (yellow) is selected in Figure 4(a) by simply double clicking on the image. Next, an additional target color is added by clicking on the green folder behind the yellow target. The snake expands to encompass both the yellow and the green targets. These images were taken with the Pyro firewire web camera. In this case, at least, little bleeding of the colors is seen. Without the current algorithm, the snake had great difficulties in crossing the yellow-green boundary. With the presented modification in the target error signal computation, the multi-color target acquisition became much more robust and simpler to use.

Conclusion

The standard statistical pressure snake routine of C. Smith was modified to function on a color image in HSV color space. If the target color is only specified in terms of hue, then large variations in target lighting and shading are permitted. This provides a significant robustness increase in the computer vision algorithm's capability to track a specific target in an unstructured, outdoor environment. One requirement here is that the target is surrounded by a frame of a clearly different hue to avoid the snake spilling over to other image components with a similar hue. By specifying the target color to contain hue, saturation and intensity values, it is possible to a establish reasonably robust method to track general image features of a common color. Here it is important to strictly penalize hue variations from the target color and only mildly penalize the saturation and intensity variations. Further, a modification to the standard pixel averaging routine is introduced. Here the target error signal is computed for each pixel of the 3x3 box first, before averaging the values. This allows us to specify the target not only in terms of a single color, but also using a list of colors. These algorithms were tested and verified by using a web camera attached to a personal computer.

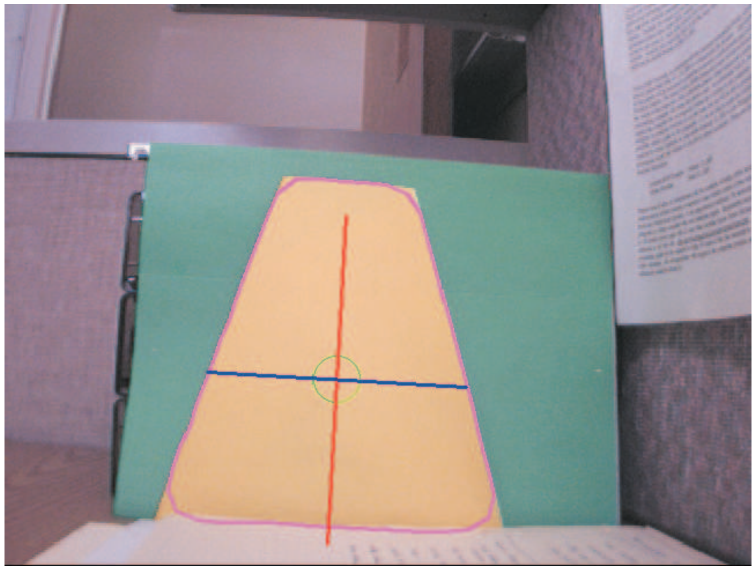

(a) Single Color Target

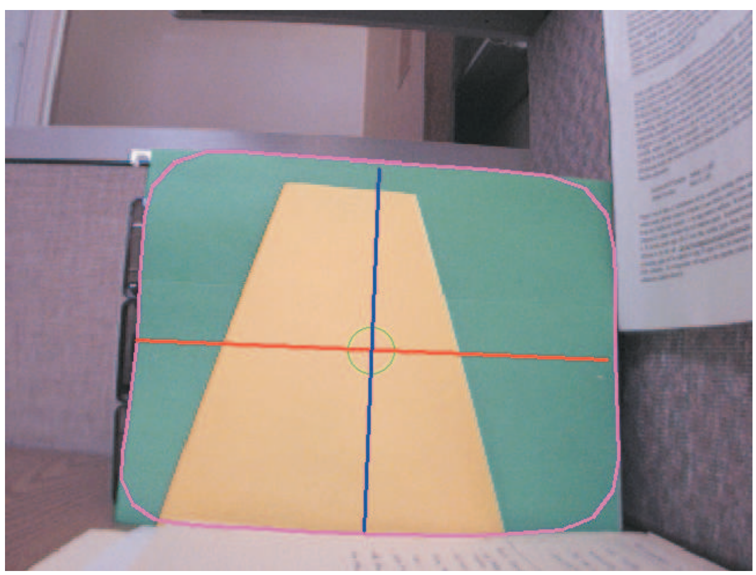

(b) Dual Color Target

Figure 4. Single and Dual Color Target Acquisitions

References

[1] Perrin, D. P. and Smith, C. E., "Rethinking Classical Internal Forces for Active Contour Models," *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition*, December 2001, pp. 615–620.

[2] Ivins, J. and Porrill, J., "Active region models for segmenting medical images," *Proceedings of the 1st International Conference on Image Processing*, Austin, TX, 1994, pp. 227–231.

DISTRIBUTION:

1 Virginia Polytechnic Institute
Attn: Hanspeter Schaub
Aerospace and Ocean Engineering Department
228 Randolph Hall
Blacksburg, VA 24061-0203

1 MS 1125
Phil Bennett, 15252

1 MS 1125
Dan Marrow, 15252

1 MS 1125
Robert J. Anderson, 15252

1 MS 9018
Central Technical Files, 8940-2

2 MS 0899
Technical Library, 4916