Using Computer Vision to Assist the Scoring of Modern Fencing

Stephanie Athow and Jeff McGough
Computational Sciences & Robotics
SD School of Mines & Technology
Rapid City, SD 57701
Stephanie.Athow@mines.sdsmt.edu

Abstract

In this paper, we describe a computer vision based scoring system to assist a referee for the sport of modern fencing. While modern fencing has three weapons, foil, epee, and saber, we focus on foil. This scoring assistant uses color blob detection to track the fencer, who has been instrumented with a color patch for the hand and a color patch for the torso. The reliability of the color sensing using different color presentation methods, various lighting conditions, and static and dynamic situations is examined and presented.

Introduction

In this paper, we describe a computer vision based scoring assistant for modern fencing. There are three weapons in modern fencing: foil, epee and saber. We will focus on foil as it is the most common of the three and requires a human to award touches. The challenges with scoring foil is that the foil must land with a point impact with sufficient force and the fencer must have priority (described below).

Determining an on-target touch is relatively easy to do with modern electronics. The foil has a button at the tip of the blade which is wired to a socket behind the hand guard. The fencer wears a conductive jacket that covers target area and a body cord that connects the foil and jacket to a scoring machine. When the button is depressed with "...more than 500 g [sic]..." [1], the scoring machine can determine if the circuit is or is not completed through the opponent's jacket and display the appropriate indicator light. To ensure enough force is applied, there is a spring inside the tip which provides the correct resistance so that the button does not depress with less force than required.

In the common event of simultaneous touches (both on-target lights display), the referee uses priority to award a touch. Priority is given to the fencer who aims at target area while extending their arm first. Priority is taken away from the attacking fencer when the opponent successfully performs a blocking action. The exchange of priority between fencers happens very quickly and can swap many times. The referee must be able to track and explain the final actions for each point, which makes it the most difficult part of scoring modern fencing.

Currently, the only available technology for determining priority when a referee is unsure or a fencer challenges a call is a slow motion replay. However the replay is still human judged, it is not widely available and it is usually reserved for the semi-final and final bouts. If this technology is unavailable, then the referee must reconstruct the actions from memory.

Our approach to priority determination is through computer vision using blob color detection. Fencers are instrumented with colored fabric patches on their hands and torso and track the centroids of those blobs. Using the relative distance between fencers hands and torsos, an expert system will compare the timing and distance changes between the two fencers. This paper will focus on the following stages: effectiveness of blob detection, motion detection, and resulting reliability of the expert system.

Fencing Background

This research revolves around determining which fencer has priority in the bout. A simplified version of the foil priority rules are:

1. A properly executed attack has priority over a counter-attack
2. A properly executed stop-hit (attack in preparation) has priority over an attack.
3. A riposte has priority over a remise (continuation of a failed attack). [2]

An attack is defined as “The simple attack, direct or indirect (cf. t.8.1), is correctly executed when the extending of the arm, the point threatening the valid target, precedes the initiation of the lunge or the fleche”.[1] We simplify this definition to be: “the hand must be moving away from the body prior to initiation of any footwork.” For this paper, a minimum threshold of 10 cm of extension will be set to be considered the start of an attack. This threshold ignores any jitter or general attack preparation motions of a fencer.

The scoring box is typically weapon selectable, as there are slightly different timings for each weapon and in epee there is no off-target. The scoring box has two sets of indication lights, on- and off-target, for each fencer. The on-target lights are colored red and green, and the off-target lights are white. For foil, there is a 300 ms lockout time, during which both fencers may hit and lights will display for both. This is a critical time for a referee to be correct in their call.

The lunge is the usually the footwork action that finishes an attack and is the action that will cause the foil to have the fastest horizontal velocity. Gutierrez-Davila et al [3] performed a study to understand the response time between an elite and medium skilled fencer. In that study, elite fencers achieved an average horizontal weapon velocity of 2.55 ± 0.42 m/s during a 601 ± 82 ms movement time to perform a lunge. A medium fencer achieved an average velocity of 1.88 ± 0.48 m/s during a 585 ± 72 ms.

Computer Vision

The CMUCam5 (Pixy) was selected as the computer vision solution to color blob detection. The Pixy is an open-source package that is a hardware and software solution. It can learn up to 7 color signatures and 7 color codes (a combination of colors) and recognize them at 50 fps. Once colors are learned, it has an on-board chip that processes the images coming from its camera and when a 3x3 or larger pixel group of a trained color is found, it will return the x-,y-location of the centroid and the height and width of the group. The 50 fps equates to 200 ms per frame, which is below the 300 ms lockout for a double touch and makes for a viable solution option to determine which fencer had priority.

The Pixy’s capabilities have been rigorously tested for reliability by focusing on color presentation, noise attributed with a static color, and dropped frames for a color in motion under various lighting conditions.

Color Presentation

Reliable color detection starts with the method of color presentation. The Pixy tracks the color's hue. Hue can be defined as "The degree to which a stimulus can be described as similar to or different from stimuli that are described as red, green, blue, and yellow." [4] Colors were presented in three manners: felt, cotton fabric and EL Wire. Training was attempted under various lighting conditions. These conditions were:

1. Fluorescents: Low (1 bank of lights on)
2. Fluorescents: Medium (2 banks of lights on)
3. Fluorescents: High (3 banks of lights on)
4. Natural Lighting Only natural sunlight through a window

Felt trained well, each color could be identified, slightly changing the lighting usually did not require retraining and no colors appeared as other colors. The cotton cloth trained less well than the felt. Orange and red were often seen as one or the other and lighting transitions usually required some colors to be retrained. EL Wire was much too dim in any lighting condition and rejected as a viable option.

Static Noise

For this paper, noise is defined as the perceived motion of the centroid location of a static color patch. In order to correctly detect an attack, this noise must be accounted for to avoid false positives. In each location, a projector screen or white wall was available to tape the color patches to and the Pixy positioned such that the color patches were roughly centered in the field of view. The Pixy was placed 96.5 cm away from the screen and moved away in 5 cm increments until to 304.5 cm. This represents the area a referee could be standing while directing a bout. A script was run for 1 second at each location to collect the color signatures and the x-y location of the centroid for each signature. Three lighting conditions were tested:

1. Fluorescents: Medium (2 banks of lights on)
2. Fluorescents: High (3 banks of lights on)
3. Natural Lighting: Only natural sunlight through a window

The low fluorescent condition was omitted because that lighting level would be too dim to safely fence.

For each location and color signature at that location, the standard deviation was calculated and histograms of the recorded x- and y-locations were plotted. An example of a histogram can be seen in Figure1. This histogram is representative of most recordings. It is a Gaussian distribution that is mean centered at the centroid of the color patch.

Compiling all the standard deviations into one plot for each lighting condition, it

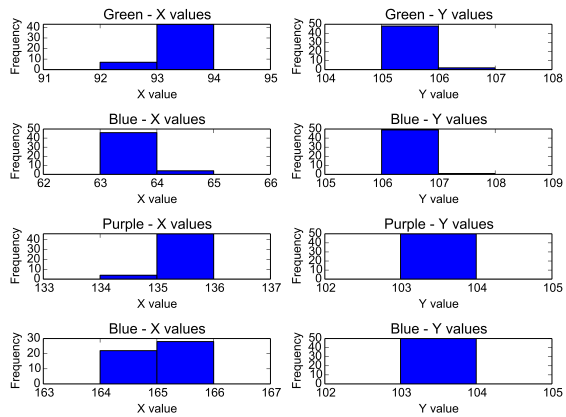

Figure 1: Histogram of Static Colors at 254 cm

can be determined where there are high standard deviations, which colors suffer from large standard deviations and how frequently the standard deviation is large. Figure 2 shows the standard deviations for all tested colors. Along the x-axis is the color name, along the y-axis is the distance location of the Pixy and the z-axis is the standard deviation value. By looking at this figure, it is easy to see that Yellow is a poor choice with its high frequency of large standard deviations. Red and Orange have larger standard deviations but aren't nearly as drastic as Yellow and the one instance of Purple. For the remaining experiments and figures, the focus is on Green, Pink, Purple, and Blue since they have low standard deviations and do not suffer from potentially being detected as another color.

Figures 3 and 4 show the results for the fluorescent medium and high cases. Green, Pink, and Blue remain at a consistently low standard deviation while Purple spikes at a low frequency.

Motion Tracking Reliability

The reliability of tracking an object can be broken down into three objectives:

1. Dropped frames and the frequency a dropped frame occurs
2. Minimum real space value of a pixel
3. Maximum bound on pixels travelled per frame

To examine the dropped frame rate, a simple stepper motor-pulley system was set up, as seen in Figure 5, and a color patch was pulled left and right at a constant velocity with no obstructions between the patch and Pixy. The motor accepted commands in RPM and three RPM values were tested: 30, 50 and 80 RPM, 80 RPMs being the maximum of the motor. Table 1 shows the conversion from RPM to linear velocity. With 400 steps per revolution and a pulley radius of 3 cm, Equations 1 and 2 are used to calculated the linear velocity of the color patch.

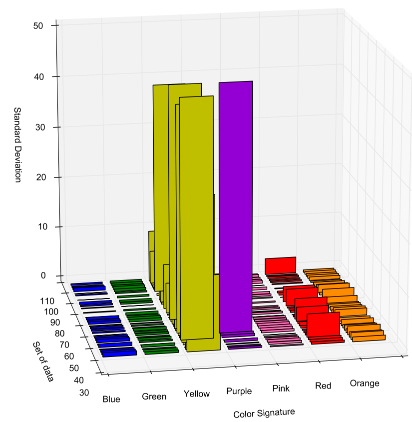

Figure 2: Natural Lighting Standard Deviations

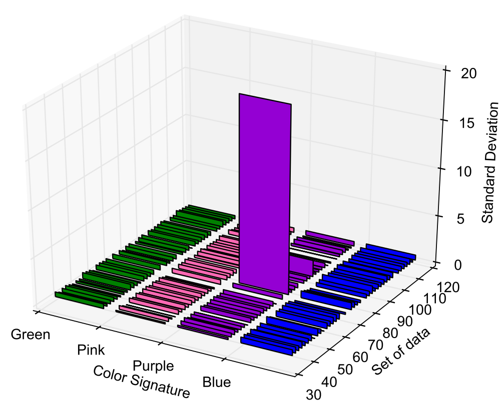

Figure 3: Fluorescent - Medium Standard Deviations

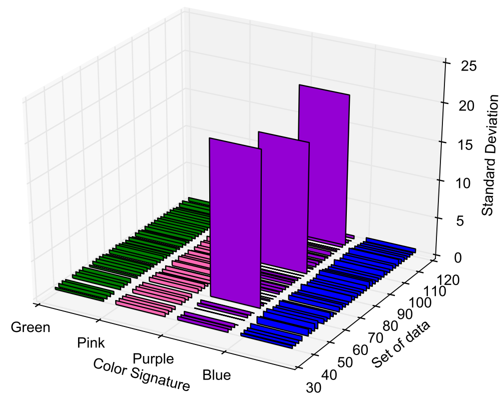

Figure 4: Fluorescent - High Standard Deviations

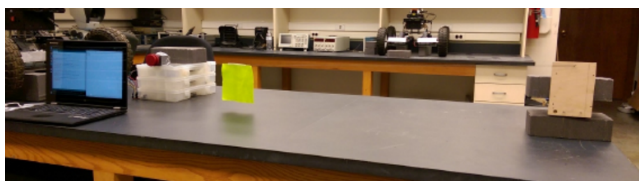

Figure 5: Motor-Pulley Set-up

$$\omega = (\theta/Step) * (Steps/s) \tag{1}$$

$$V = \omega * r \tag{2}$$

The Pixy was positioned to maintain the color patch within the field of view for the full range of the system. It was positioned at 96.5 cm away from the center of the system and moved in 5 cm increments to 182.5 cm. At each location, the color signature was tracked left and right at 30, 50, and 80 RPMs. When the Pixy detects a color, it returns the frame number and signature information for that frame. Using this information, each file was checked for any missing frames by counting the number of times the color signature was recorded against the final frame number. These were

RPM	Linear Velocity (cm/s)
30	9.4
50	15.7
80	25.1

Table 1: RPM to Linear Velocity of Color Patch

Distance to Object (m)	Pixel Width (cm)	Pixel Height (cm)
1.0	0.24	0.22
1.5	0.36	0.33
2.0	0.48	0.43
2.5	0.60	0.54
3.0	0.72	0.65

Table 2: Pixel Space to Real Space at Various Distances

summed across all files. The dropped frames is the number of counted signatures subtracted from the total frames reported back from Pixy. The final result was zero dropped frames.

To find the real world value of a pixel, Pixy has a known 640 x 400 resolution with a 75° field of view in the x-plane and a 47° field of view in the y-plane. Table 2 displays the conversion from pixel space to real space based on the distance between Pixy and an object. The maximum value of a pixel width is 0.72 cm, which provides enough granularity to detect the 10 cm minimum extension to start an attack.

To calculate the upper bound on pixel change between sequential frames, lunging speeds are known for an elite and medium skilled fencer. The elite fencer has an average of 5.1 cm/frame and a potential of 6.2 cm/frame during a lunge. A medium skilled fencer has an average of 3.8 cm/frame and a potential of 4.7 cm/frame during a lunge. Taking 6.2 cm/frame at 1 meter, an elite fencer can travel up to 26 pixels per frame. At 3 meters, an elite fencer can travel up to 9 pixels per frame.

Conclusions

Felt initially appeared to train very well. Considerations such as wool allergies and fencer comfort led to choosing the cotton cloth, which trained slightly less well than felt. By looking at the noise histograms and the standard deviation bar graphs, all colors have Gaussian noise and some colors train significantly better than others. It can also be noted that when a color has a low frequency of high standard deviations, the high standard deviation is much larger than the rest of the deviations.

Applying reasonable lower and upper bounds to pixel changes of an object is feasible. In the space a referee would stand away from the fencers, the real space pixel dimensions are significantly below the defined 10 cm of extension that constitutes an attack. By calculating the maximum distance a fencer could travel in a frame, an upper bound on pixel change between sequential frames can be applied. This upper bound can be used to ignore the seemingly random high standard deviation.

These results show a system to track a fencer is viable using this technology. Cotton cloth provides color hues that the Pixy can reliably sense under three lighting conditions. The noise of the Pixy has been characterized and the maximum pixel change between two sequential frames has been calculated so outlier readings can be rejected in real time.

Using this technology we can determine the motion of fencers. It will serve as the foundation for the expert system that will determine priority and award the touch to the fencer with priority.

Future Work

This paper lays the foundation for future work in designing a robust algorithm to determine priority based on the hand moving away from the body using the Pixy. This algorithm would need to train the colors on a fencer standing still so that the noise levels could be adjusted accordingly. Additionally, the legs of the fencer could be instrumented with one color code per leg. By tracking the legs, a crossover event may be included in the algorithm.

References

[1] United States Fencing Association, Inc., *Fencing Rules*, September 2015.

[2] *USA Fencing Refereeing Handbook*, 1.2 ed., 2012.

[3] M. Gutierrez-Davila, F. J. Rojas, R. Antonio, and E. Navarro, "Response timing in the lunge and target change in elite versus medium-level fencers," *European journal of sport science*, vol. 13, no. 4, pp. 364–371, 2013.

[4] M. Fairchild, "Color appearance models: Ciecam02 and beyond," in *tutorial notes, IS&T/SID 12th Color Imaging Conference*, 2004.