Basketball Player Identification by Jersey and Number Recognition

Jeremy Ephrati
Department of Electrical Engineering
Stanford University
Stanford, CA
jephrati@stanford.edu

Navid Moghadam
Department of Electrical Engineering
Stanford University
Stanford, CA
navidm@stanford.edu

***Abstract*— Recognition of players in pictures of sporting events is an approachable but tough task. In the case of an NBA game, this task can be accomplished by compartmentalizing the job. By characterizing jersey color with MAP detection, isolating the jerseys and numbers, and using template matching, we can make a max-effort algorithm that tried to identify as many players as possible.**

Keywords-jersey color, in-game picture, RBG color space, template matching

I. Introduction

Image Processing is a valuable tool that is applied on media presenting sports frequently. One well known system that is used frequently is the 1st & Ten system used to generate yellow lines symbolizing the 1st down barrier in football in real time during a game. The system is also utilized for applications like advertisements in unused sub-windows, and depicting race car position for NASCAR races.

Player Identification is an application of image processing that is a topic of interest. The idea to find and recognize players' numbers was explored in [1]. The authors manipulated the HSV characteristics of the team jersey to isolate them. Then the characteristics of the image and sport are manipulated to isolate the number for image processing and identification.

In [2]. we see a different application trying to identify digits in an image: the authors try to identify cars by reading the license plates. Morphological Image Processing is incorporated to isolate the numbers from the rest of the image, and the statistical method of correlation is used to detect the digits with an increased recognition success rate over [1].

Previous documentation on color properties provides motivation for the framework of this paper. In the following document, [3], the author investigates the viability of using color detection in MATLAB to aid the visually impaired. He mentions the idea of recognizing certain colors, which inspired the idea to recognize jersey color by RGB value.

In this paper, we investigate the viability of a player detection system that has knowledge of the visual characteristics of an in-game image. We will use MAP detection to isolate jersey regions in the image, image segmentation to further process the image, and an OCR based number detection method to guess the number of a player. The input to our algorithm is an image and knowledge of the team and jersey type, and the output will be the input image with the names of the players.

II. Jersey Recognition

Since we will know the teams that will be playing in the picture, we will have up to 4 possible jersey colors present at the games (each team having a home and away jersey). In basketball, since the jerseys are typically one color without any stripes or advertisements, we can characterize them as a certain color. We attempted finding a loose range of RGB values that would characterize the jerseys in the pictures. The proposed algorithm would go through every pixel, and check whether its RGB values fell into the respective acceptable ranges. If all three of its RGB values did not fall in the range of acceptable values, then the pixel got assigned RGB values of 0, to denote that it was not a jersey. Unfortunately, this algorithm was far too simplistic to succeed. Even with the removal of the crowd and its wide range of colors, there were still some false positives associated with this method. Additionally, jerseys did not always get identified by the algorithm. The algorithm was not robust for the different varieties of pictures that could possibly get taken, because of variables like angle, brightness, cropping, etc.

To circumnavigate that issue, we looked into designing MAP detectors for every jersey. We found five training images per jersey for our MAP detector. We attempted to find jerseys in different lighting conditions to fully capture a wide range of RGB values for the maximum chance of identifying a jersey. The regions with the color of interest were painstakingly cropped. Then, we train our MAP detector using our training images and derived masks. To isolate jerseys in an image, we will apply our MAP detector to the image. This algorithm ends up being much more robust than the previous simplistic one.

Figure 1. Result of MAP detection

Ideally, this will produce a mask of the image, with the jerseys being white and the rest black, as shown in Figure 1. Realistically, false positives are impossible to avoid. Therefore, we will have to do more processing to obtain our jerseys without any noise from falsely identified jersey pixels. We can use the regional properties of the noise to discriminate against them. Each basketball team is allowed five players on the court. With that logic, we will find the regions with the five biggest areas, and eliminate those that do not fit that criterion. This will leave us with some noise, and our jerseys of interest.

III. Number recognition

Due to the steps we have taken previously, at this point we have a mask of the jerseys and some additional noise. We want to make sure that the numbers are isolated from the rest of the image, with no connected edges. We apply an erosion with a small disk structuring element to create some separation between the number and other elements of the picture.

To recognize the player, we need to try to recognize the number on the jersey also. We will look at each player individually, and compute which player it has the greatest chance of being. To isolate one player, we loop through the centroids in order of decreasing area. We look at each “region,” or jersey by itself, but setting the rest of the regions to 0, and arrive at a mask like that in Figure 2. Then we look to isolate the number of the player. To do that, we invert the mask, and then remove the largest area, which will be the background. We will be left with the number of the player and extraneous details from the jersey, which is usually the Lakers logo or the player’s name, depending on whether the player is facing forwards or backwards during the photo.

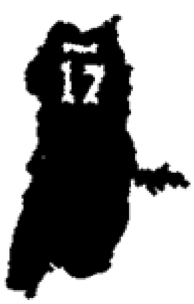

Figure 2. Result after additional processing

Now we have what we need to recognize the player. Our first idea for doing so is to use template matching. We constructed five templates, one for every possible player. We then found the height of the number on the jersey. Using this, we resized each of the templates to match the height of the number. After resizing the templates, we applied a convolution on the mask of the number with each of the templates. We then observed the maxima of the convolution results to see which template had the highest chance of being a match. However, many factors deteriorate the viability of this strategy. One factor that seems to affect the accuracy of template matching is the angle of the number. The slight tilt causes a lack of alignment that will increase the chance of a false positive. Also, a jersey can fold over itself, which will cause distortion of the number.

Due to the inaccuracy of the template matching method, we decided to try OCR to detect the number. The OCR we used was designed to read a text file of numbers and letters. It analyzes each character and computes a correlation with every other template, and picks the best match. We edited the source code to only consider numbers as the possible results. However, testing revealed some flaws. Very frequently, characters would get confused for other characters. An example of a pair of numbers that would commonly get confused for each other is ‘1’ and ‘4’. However, for this application, we do not need an OCR engine that can recognize any number. We only need it to distinguish between five possible numbers. Restructuring our identification method to only account for those five possibilities will cut down on errors during number recognition.

Therefore, for two digit numbers, we will use two parallel OCRs to detect the jersey number. Each OCR will be customized for each digit. We know that the Lakers have five numbers on their starting roster: ‘7,’ ’15, ’16, ’17’, and ‘24’. Therefore, we know the first digit of the number will either be 1 or 2, and the second digit will be a 4, 5, 6, or 7. Using this fact, we can run two parallel OCRs. We will construct one OCR for the first digit (that only recognizes 1 and 2), and another OCR for the second digit (which only recognizes 4, 5, 6, and 7). We then take the results of both OCRs and combine the digits to complete the estimate of what the number is.

The analysis of a single digit number is a simple derivative of the previous analysis. We will just use the OCR for the second digit of the two digit numbers to find the number.

After getting the results of our OCR number recognition, we check for any exact matches. If we get a number that does not match, we will use some logic specific to the Lakers team. Since we know that there is only one player with a single digit number, we can claim that any single digit number corresponds to that player. The other four numbers are 24, 15, 16, and 17. Something we can exploit from these numbers is that their second digit is unique. Therefore, we can depend solely on the second digit to find what the number is.

We run the risk of incorrectly identifying players when there are less than five on the court. Therefore, we must do some additional processing to remove extraneous areas. We can characterize jerseys as having regions inside that have a larger than unity height to width ratio, which are numbers. We also know numbers typically take up a roughly constant portion of a total image. Therefore, we can first eliminate any noise inside the jerseys that are smaller than a certain percentage of the total image, such as .02% like this method [1]. Then we find the bounding boxes of all the remaining regions, and only keep regions that have a height to width ratio greater than the threshold, which will be something greater than 1. This will remove spotty noise. An additional result of this processing is that the logo of the team, which hovers above the number of the player, will be removed.

If all goes according to plan, only players will be processed by the algorithm. After we have our guesses for each player's identity, we can print the players' names right under their numbers to demonstrate to the user the results, as shown in Figure 3.

Figure 3. Result of Algorithm.

IV. RESULTS

We tested our algorithm using 30 in-game images with varying sizes and characteristics. Our figure of merit was the number of correctly identified players over the number of players total. We chose to not include pictures where the number was obstructed or altered, because of the algorithm's vulnerability to such features. Our testing produced an accuracy of 53.6%.

Figure 4. Example of a figure caption. *(figure caption)*

V. ANALYSIS

Though there was great care in trying to make a robust algorithm to detect players, there were m0any factors that reduced the efficacy of the algorithm. One issue with the algorithm is that it relies on a faulty method of detecting and identifying text. Frequently, the OCR engine would misidentify a number, which would lead to a wrong identification of a player. The OCR engine is vulnerable to any imperfections in the depiction of the number in its relation to the template. For example, if the number being detected is rotated in any way, the results of the OCR detection will frequently be incorrect. Another problem with OCR is that if the number being detected is too thick. This can cause the probability of false positives to increase significantly. A possible result of the mentioned vulnerability is shown in Figure 4.

Figure 5. OCR Failure leads to wrong player getting recognized

Another issue that arose was that many pictures had a lot of noise with similar RGB properties to the Lakers' jerseys. The MAP detector would detect large regions of noise as Laker's players, and the processing between the MAP detection and OCR detection did not eliminate the noise. This sequence of events potentially leads up to something akin to Figure 6, where a "phantom" player is identified. The prevalence of this error can be prevented by more aggressive processing to remove noise. However, the tradeoff to more aggressive processing can lead to more errors in identifying the number, due to distorting the jersey or number. The best solution might be to design a more robust MAP detector, to exclude more noise and reduce the chance of a false positive jersey getting recognized.

Figure 6. Phantom Player recognized

Another issue that arises is the nature of the input picture. Many characteristics of the input images can impact the efficacy of the algorithm. The shading on the jersey can impact the output of the map detector, and cause issues if there is distortion of the MAP detector output. Also, if someone's arm cuts across the jersey or the number, it can cause issues by either ruining the isolation of the number or splitting the region of the jersey. Ripples in the jersey also can cause issues, because they can overlap the number on itself and cause errors in the OCR detection. All of these errors can lead to no player getting recognized in the image, as shown below in Figure 7.

Figure 7. No Player recognized

VI. CONCLUSION

The goal of this project was to investigate the viability of constructing a detection method for identifying players from the NBA in an in-game image. Though our focus was scaled down to one team and one of their jerseys, this does not take away from the results of the paper. We found limited success with our algorithm, but there is a lot that can be done to increase the robustness of this algorithm. For example, the method implemented in [3] can be used to improve the text recognition for the jersey number. Overall, with more robust individual components to our algorithm, and more time to incorporate other teams into the system, there is great potential for a system that will identify players from any team.

ACKNOWLEDGMENT

This project was completed with the tools Jeremy and Navid learned from the course EE 368: Digital Image Processing. Jeremy and Navid thank David Chen and Andre Filgueiras de Araujo for their help in guiding them in the project and instilling a conceptual understanding of the subject matter at hand. Jeremy and Navid also thank Bernd Girod for teaching them the tools essential for this project.

REFERENCES

[1] M Saric, H Dujmic, V Papic, N Rozic, “Player number localization and recognition in soccer video using HSV color space and internal contours,” World Academy of Science, Engineering and Technology 43 2008

[2] S Ozbay, E. Ereelebi, “Automatic Vehicle Identification by Plate Recognition,” World Academy of Science, Engineering and Technology 9 2005

[3] I. Z. B. A. Aziz, “Object’s RGB Color Extraction”, Nov 2007, unpublished

Jeremy Ephrati implemented the algorithm devised by coding it on Matlab; he also constructed the poster. Navid Moghadam conducted testing for the algorithm and wrote the paper.