Identification methods For Robotic Systems

Two approaches for computers to identify symbols on Playing Cards

Edward Giardina, University of Massachusetts Lowell
http://www.edgiardina.com/projects

ABSTRACT
Working with real-life objects and integrating their properties into systems that robots and other embedded systems can identify is a non-trivial task. The fundamental ways for these systems to be read by robots and other systems has been either to modify the world for easier reading for the robot, or to use the existing properties of the object and have the robot decode those properties. In my blackjack playing robot, I explored both of these systems in order to better understand the interaction that robotics must make in the real world to do things like play games. First using barcodes, then using cameras, I explored different ways that robots could identify standard playing cards.

INTRODUCTION
The playing card is a mediaeval artifact: it has been with us for over 600 years. There is a fascination in the designs, the tactile pleasure, neat symmetries and quirky symbols.[i] Millions of dollars are exchanged daily on the simple overturn of a stiff piece of paper with an austere picture on it. And within the last fifty years, the fate of these cards has been tracked electronically, from mathematicians like Edward Thorpe, to others hoping to gain financial independence from casinos. With the advent of blackjack card counting systems and other statistical analysis of how playing cards work, it is of some importance to analyze how independent computing systems might read playing cards in real life scenarios.

The problem is, playing cards are more difficult to read for machines, since machines don't read inherent cues from reality like humans do. Humans can easily and quickly orient and image in their heads, so a sideways Six Card is still a Six Card. The inherent way for people to identify playing cards is visually. However, vision systems on robots are still a relatively new field. This paper explores both Barcodes and Cameras to identify these playing cards.

BARCODES
The identification and analysis of objects in a real space is a long-standing problem in the field of robotics. It is also a major concern in the commercial space of retail, as well. Before graduate students can build robots to analyze playing cards, more pressing issues like analyzing products in a supermarket were tackled. Barcodes, made popular by retail and supermarket chains, have become the lifeline of retail stores inventory and Point of Sale systems. Being able to not only easily identify products, but also track their movement in and out of stores, meant huge savings for these stores. Even though objects sold in stores range in

size, color, texture and shape, barcodes help universalize them for easy electronic identification. The playing cards I am working with are actually, from an objective standpoint, not as difficult to identify. Each card has a set dimension, material and the symbology. However, even uniform items can be difficult to identify. While humans can easily pick out what makes a Jack a Jack, computers are very literal, and things like orientation, lighting, and other factors play a role in what an object appears to be. By using barcodes on playing cards, like I did in my blackjack robot project, I eliminate two major difficulties in reading: orientation and light levels. Barcodes can be read either upside down or right-side up, just like the symbols on playing cards themselves.

Playing cards with barcodes affixed

While I only affixed barcodes to one side of the cards, as shown in the above image, I could have went further. The barcode could be printed on the top and bottom, like the numeric symbol found in each corner. This would allow quick reading regardless of where the barcode was located. The card would have to pass only halfway under the scanner regardless of up-down orientation.

Barcodes worked well under multiple trials. However, I did encounter situations in which the barcodes didn't read. I partially attributed this to the fact that I taped the barcodes to the playing cards, and the tape caused inadvertent reflection on the surface of the cards. Since barcode readers use reflected light, it is quite possible this was the source of some of the failed reads. The other issues encountered were those of distance and angle. The cards have to pass their barcode under the reader at a certain distance and angle. Depending on the rules of certain casinos and of certain card games, this might be impractical. For instance, in blackjack tables at Foxwoods casino in Connecticut, players are not allowed to touch the cards. This is simply to prevent cheating, but the distance that a barcode reader would need to be is within three inches, too close for most gaming establishment's comfort.

The real problem is that no playing card company makes barcoded cards. No casinos would ever modify their casinos to handle such cards. The real-life implications are that modifying the environment may work in a small trial, but in a system used worldwide with a fair amount of regulation, it just isn't practical, especially in something as pervasive as playing cards. In creating my project, I knew this to be true. However, I wanted not only to have experience using easily readable codes like barcodes, but I also wanted to know how possible it is to count and maintain card status using microprocessors.

VISION SYSTEMS

Of course, one can note that playing cards are specifically unique, and that modifying them isn't necessary to determine its value. Indeed, the visual cues that determine a card's value to a person can also be ascertained by machine, as well.

For my project, I wanted to use a low-cost, easy to implement camera system. I'm sure that very expensive camera and software packages could do what I want to do, (identify a playing card's value) but I want to keep both the interface simple, and the ability to interface directly with a microcontroller available. My needs were relatively simple: a low-cost, moderate performance digital camera that could produce decent resolution and had little problems shooting items close-up. The CMUcam from Carnegie Mellon University seemed like an excellent fit. The camera had direct TTL and PC serial interface, was relatively cheap, and had several image analysis functions built in.

the CMUcam

Identifying a card's unique one character value indicator would be perhaps the easiest way for a person or robotic system to identify a card. All playing cards have either a number or letter representing its value in the opposite corners of the card. Software that performs this kind of recognition is commercially available, and it is called OCR, or Optical Character Recognition. It is possible a still image from a CMUcam or like camera could be fed into an OCR program, which would report back characters seen on a card. However, modern OCR systems are geared towards scanning solutions, in which color matters very little, and orientation and position are negligible factors.

Another possibility for identifying the value of a card lies not in interpreting the symbols on the card itself, but rather taking a 'hash' of the color and light data given off by the camera's image of the card. The CMUcam has several built in values that it reports back to the user on command. Color data, light values, mean color values, standard deviations of color values, all these are available from images viewed on the CMUcam. And since each playing card has a different appearance, its reasonable to assume, given identical ambient light levels, that each card would contain a slightly different value for these measurements of color, light, etc.

However, this approach assumes a few caveats. The first is that to measure that color or light data, the card would have to be placed in the same position and read identically each time. This isn't exactly ideal. The whole concept between moving between the barcodes and the camera was that the environment is supposed to be less-modified. What I found when conducting preliminary readings of light is that the camera is very sensitive. I found myself constructing a small viewing window so that the camera would only observe light that I allowed it to, helping to eliminate fluorescent flickering and shadows from observers and other equipment.

Oftentimes, in a controlled light situation, (like using an external lamp and placing the camera in a cup or box, preventing shadowing) , where the camera wasn't moved, the color values fluctuated as much as three integers up and down (where a color value can be between 10 and 240). A certain amount of error is to be expected, but the margin of error on color values and color standard deviation was greater than some of the differences in values between cards. Therefore, the average color value of red in the number five may be 120 (on a scale from 10 to 240) with a fluctuation of +/- 2 units. However, the number six may have a color value of 121, with the same fluctuation. Therefore, a color value of 120 may indicate a five-card's average color, or a six-card's average color with a small error.

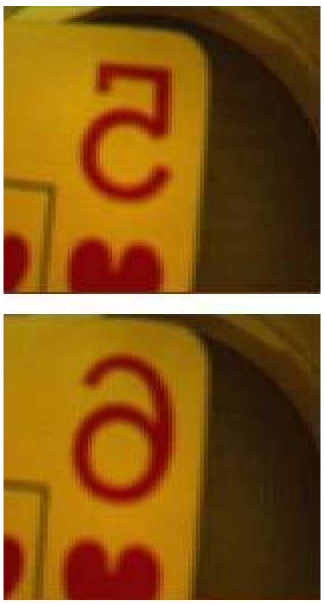

A five card and six card, taken from the CMUcam's frame dump feature.

The method of measuring color met with limited success in my dealings. The difference between the color of an Ace and the color of a King was indeed noticeable in comparing values. (For the purpose of consistency and proof of concept, I only measured red heart cards) However, in attempting to determine the difference from closely aligned numbers, like five and six (see above image), the camera faired poorly.

In theory, with infinitely sensitive and precise equipment (and perhaps infinite funds) the amount of red ink atoms or red ink in picoliters could be measured and determined. (whether printing presses could lay down the ink that consistency is also of doubt) The idea is that a card isn't just the symbols displayed on it, it is also a certain value of reflected light, and it is also a certain color value. Using devices to measure these values instead of using complex algorithms to identify symbols was the goal of the CMUcam trials I ran. Complex algorithms are less than elegant on small, microprocessor based systems. This was an attempt to simplify something complex like image analysis with an elegant solution.

Other reasons why this solution is elegant is that it would require very little memory to operate. Theoretically, one integer value between 0 and 255 would be the only item stored for each card value. Instead of taking photos of playing cards and making an active comparison, which would take more time and cost more memory, this approach needs only to know approximate values of each card based on a histogram taken at some time.

But the camera used on the CMUcam as well as the technology used on almost all digital imaging systems is one meant to be pleasing to the human eye, and not necessarily precise methods of light measurement. The camera uses lots of tricks to assemble an image based on the light it receives. The camera is also incredibly low resolution.

The camera is also a CMOS based camera, and not a CCD based camera. CMOS cameras are inexpensive and require little power, which makes them ideal for microprocessor applications. However, CCD cameras provide high-quality, low-noise images.[ii] Perhaps with a CCD camera, random noise might be less present, causing a smaller margin of error in color values.

To attempt to keep the light values the same for each trials, I took an upside down plastic yogurt cup and sliced a slot into which I could place a card. I cut out the top so that the camera could be placed atop it. This allowed me to have only a small hole from which light could pass, and it meant I could control the size of that hole to maintain consistent light levels. This also coincided with my attempt to keep the cards lined up at the same position for each read.

Barcodes Compared to Cameras

However, this modification of the environment was exactly the type of activity I was trying to prevent in this experiment. The barcodes at least, while the cards themselves were modified, had a certain flexibility in use. The barcodes could be read at varying angles and were not up-down orientation specific. The camera experiment meant unmodified cards, but it meant the cards had to be in a specific position the entire time. Outside of a lab, this wouldn't be practical at all.

Also advantageous about barcodes is that the technology is very pervasive and therefore rather refined. Supermarkets around the world use them, and since it's an economic necessity to have a low non-read rate, advances to improve read rates have been made. For instance, several barcode readers exist with 'star pattern' lasers. That way, orientation is almost a non-issue.

Camera technologies are well defined, but finding the right camera for this application may be difficult. The CMUcam has built in bounding, image tracking, and all sorts of useful mobile robot applications. However, it does not have built in algorithms for image identification. The CMUcam is programmable, and perhaps given more time and a broader exploration, the CMUcam could be written specifically for identifying playing cards.

Future Possibilities

Even more exciting is the move away from laser based barcodes. CCD barcodes, which use digital camera-style reading, are becoming popular as well. **"Why a CCD scanner?** Compared to laser scanners, CCD barcode scanners are extremely durable and require less maintenance. CCD scanners consume little power and work under most lighting conditions. The CCD barcode scanners … all have built-in decoders and can auto-discriminate several bar code symbologies."[iii] This means no laser light, which means less obtrusion into the environment. Casinos would be quick to notice laser light peeking out of someone's sleeve, after all.

This development leads me to believe a convergence may be possible between barcode readers and CCD based systems. If barcodes can be identified using camera-like systems, why can't playing card numbers be attacked the same way? After all, both barcodes and playing card symbols don't change. Since they don't change, there is a limit of how many different symbols we could possibly see, which means a computer or robot system could examine a certain card or barcode or any other symbol, take some sort of hash of the data, and compare it to the limited symbols we know exist in the data set.

While these systems wouldn't use such a simple algorithm as the one I proposed, they most certainly would have to be microprocessor based.

Conclusion

Unfortunately, my experiments failed to produce the data needed to create the Camera-using blackjack robot. However, I did manage to gather quite a bit of data for the idea, and I believe it could be created given either a rewrite of the CMUcam software, or perhaps a different means of visualizing the cards. The blackjack robot described in this experiment can be found at http://www.edgiardina.com/projects/

I will also be posting CMUcam shots there, as well as this paper.

---

[i] The World of Playing Cards Website, http://www.wopc.co.uk/ (A brief playing card history with global card comparisons)

[ii] Howstuffworks website, http://electronics.howstuffworks.com/digital-camera3.htm

[iii]ID automation's product website http://www.idautomation.com/ccdreaders/