2007-01-01

Using Computer Vision to Create a 3D Representation of a Snooker Table for Televised Competition Broadcasting

Hao Guo
Technological University Dublin

Brian Mac Namee
Technological University Dublin, brian.macnamee@tudublin.ie

Recommended Citation

Hao, G., & MacNamee, B. (2007). Using Computer Vision to Create a 3D Representation of a Snooker Table for Televised Competition Broadcasting. *AICS:The 18th Irish Conference on Artifical Intelligence & Cognitive Science,*UCC, Cork, 29-31 August. doi:10.21427/66rq-jj82

Using Computer Vision to Create a 3D Representation of a Snooker Table for Televised Competition Broadcasting

Hao Guo and Brian Mac Namee

School of Computing, Dublin Institute of Technology, Kevin St., Dublin 7, Ireland
{Brian.MacNamee@comp.dit.ie}

Abstract. The Snooker Extraction and 3D Builder (SE3DB) is designed to be used as a viewer aid in televised snooker broadcasting. Using a single camera positioned over a snooker table, the system creates a virtual 3D model of the table which can be used to allow audiences view the table from any angle. This would be particularly useful in allowing viewers to determine if particular shots are possible or not. This paper will describe the design, development and evaluation of this system. Particular focus in the paper will be given to the techniques used to recognise and locate the balls on the table.

1. Introduction

When watching televised snooker competitions it is often hard to see whether it is possible for a player to make a particular shot unless we can see the table from the player's point of view. Unfortunately, television cameras cannot always disturb the player to get the correct viewing angle. This paper will describe the Snooker Extraction and 3D Builder (SE3DB) a system developed to generate a 3D model of a snooker table from an overhead image of the table. This 3D model could be used in television broadcasting to allow viewers view the layout of the balls from any angle and so determine whether or not the tough shot the player is eyeing up is really possible.

Technology has been creeping into sports broadcasting more and more over the past number of years, and the following section will describe some notable examples of this to serve as a background for the discussion of our system. Next, the system itself will be described, including an overview of the digital image processing techniques utilized to extract the positions of the balls and a particular observation which makes this possible.

Finally, we will describe an evaluation that has been carried out on the system and suggest the directions in which we expect the work to go from here.

2. Background

Computer vision, coupled with technologies such as augmented reality (AR), has high potential for enabling a new class of application in television broadcasting [6]. Effects such as highlighting particular players, displaying team logos and illustrating distances have now become commonplace. One of the earliest examples of AR being used in sports broadcasting is the FoxTrax system [1]. This system was developed to highlight the location of the puck in televised ice hockey games. As the puck can be extremely hard for viewers to see when moving at high speeds across the ice, Foxtrax added a virtual glow and a comet trail to it. To achieve this, the developers of Foxtrax created an instrumented puck which emitted infra red pulses which were captured by cameras positioned around the ice rink. This system used 10 cameras, each of which focussed on a particular portion of the field, meaning that during a match, the puck was always in view of at least one camera. While the Foxtrax system was highly successful, drawbacks include the large amount of hardware required and the fact that modifications had to be made game equipment itself, which is not always possible.

The Hawk-eye system **(**www.hawkeyeinnovations.co.uk**)** [5] is another example of vision-based technology used in sports broadcasting. Originally developed for use in cricket matches, the Hawk-eye system creates a 3D virtual simulation of a sporting event which can then be played back and viewed from any angle to review the action. Since its original deployment Hawk-eye has gone on to be used in tennis and, more recently, snooker. Interestingly, in tennis Hawk-eye is not only used in broadcasting, but also by match officials to review calls on whether a ball is in or out.

Hawk-eye is particularly interesting in relation to the work being presented in this paper. Not only does Hawk-eye have similar goals to our own work (to create a virtual representation of a sports event) but it also focuses on the same sport (snooker). Unfortunately, as Hawk-eye is a commercial product there is very little information available in the literature about how it works.

As augmented reality has been used more and more in broadcasting, people have become more demanding in terms of quantity and quality of the visual information. Also, viewers want to interact with the content or influence the presented material on

TV. PISTE [2] is a system aimed at addressing all of these needs, providing broadcasters with the tools necessary to create enhanced content at transmission time, and the viewers with set-top-box technology capable of handing requests for interaction personalization. This system is capable of performing measurements, frozen shot display and comparison of different attempts displayed simultaneously.

3. Development

The purpose of this system is to extract information about the contents of a snooker table, and from this information build a virtual model of the table. Snooker is a particularly attractive sport for this kind of application as the contents of the environment are entirely controlled, and a large amount of information is fixed – e.g. the relative dimensions of the table and the balls. The development of our system involved the creation of a series of prototypes which attempted to create a virtual 3D model of a snooker table. This task can be divided into identifying the boundaries of the table itself and identifying the positions and colours of the balls on the table. The development of our system will be described in these two parts. Our system assumes that a human operator determines when a picture of the table should be taken and a subsequent 3D model created. The operator should only do this when the image of the table is clear of players, snooker cues, rests etc, and so we do not consider the presence of any of these objects.

In developing this system OpenCV (sourceforge.net/projects/opencvlibrary) was used for most of the computer vision algorithms used and OpenGL (www.opengl.org) was used for the display of the virtual table.

Extracting Table Boundaries

Detecting the boundaries of a snooker table is a relatively straightforward task as so much is known about what we expect to see in a typical snooker scene. In our system the boundaries of the table are extracted through edge detection using the Canny edge detector [7] followed by the Hough transform [3]. The results of a Hough transform is a set of candidate edges for the table boundary. From this candidate set of edges a representation of the table boundary can be extracted by using domain information about how we expect a table to appear. Figure 3 shows a series of images which illustrate the steps required in identifying the table boundaries from an image of a snooker table.

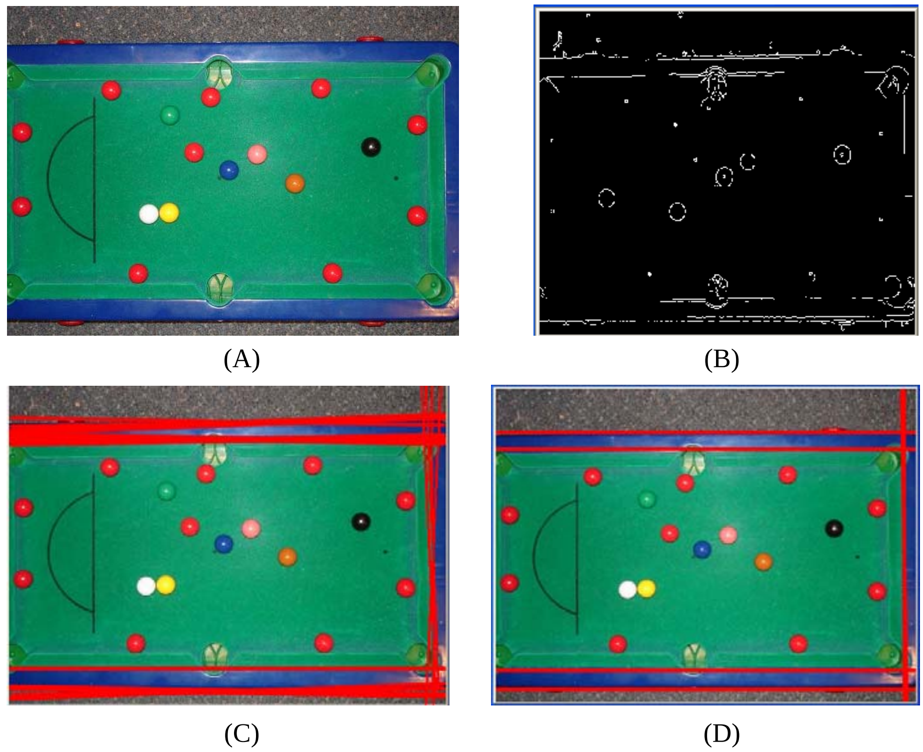

Fig. 1. (A) The original image, (B) the edge detection result, (C) the results of a Hough transform super-imposed on the original image and (D) the final extracted table boundary super-imposed on the original image.

Identifying Snooker Balls

The first prototype developed to identify the balls on the snooker table used a very simple approach which performed an edge detection on the image of the table and used a circle-based Hough transform to extract the outlines of the balls. Unfortunately, the Hough transform did not perform particularly well. It is believed that the reasons for this are that the image was too noisy and that the circles extracted from the edges of the snooker balls were too far from perfect circles which caused problems for the OpenCV implementation of the Hough transform used.

The second prototype developed used a little more of the domain information available to us. This time it was realised that since the baize of the table was known to be green it was obvious that the balls should stand out prominently from it. Hence, by using a simple flood fill algorithm [3, 4] the ball positions could be extracted without having to rely on edge detection. From this result, connected components can be identified as the balls on the table. Unfortunately, while some success was achieved using this technique, this prototype suffered from the fact that a number of the balls (in particular the blue and green balls) were particularly close to the colour of the baize and so were not identified successfully.

While developing the second prototype the key observation which resulted in the success of the final system was made. It was noticed that there is a strong specular reflection (a very bright spot) on the surface of each snooker ball. An example of this is shown in figure 2. This appears reliably in images of snooker tables, including those broadcast on TV. It was decided to take advantage of these bright spots to detect the balls using simple thresholding [3]. By converting the captured image of a snooker table into a grey-level image it is a straightforward task to set a threshold value which will extract only specular reflections (plus occasionally the white and yellow balls).

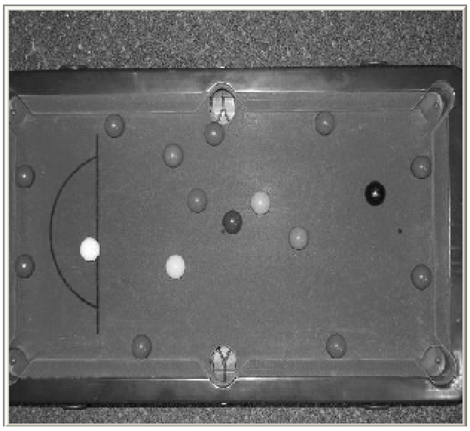

Fig. 2. A greyscale image of a snooker table which shows an obvious specular reflection on the surface of each ball.

Figure 3 shows an image of a snooker table, the results of thresholding and the final identified snooker balls. From figure 3 (B) it can be seen that, even though most of them are just tiny points, the specular reflections on the surface of each ball are particularly obvious. After the image is thresholded connected components analysis is used to extract each specular reflection, the centroid of which is considered to be the location of the ball. The white and yellow balls appear as particularly large objects in the thresholded image. This is not a problem, however, as the centroid of these objects still gives a good position for the ball.

(A)

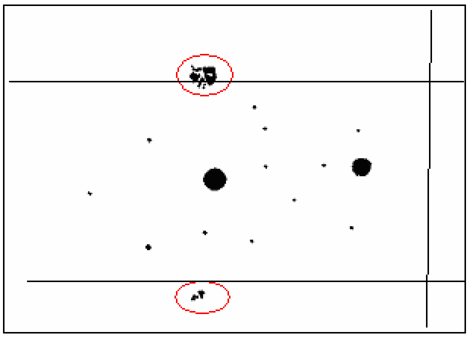

(B)

(C)

Fig. 3. (A) The original image. (B) The results of thresholding. (C) The final set of detected balls highlighted with green dots.

From figure 3 (B) it can be seen that there are some spurious results in the thresholded image around the pockets of the table. These are the result of specular reflections arising from the material used to coat the insides of the pockets. These results were omitted from the final ball identification result (as shown in figure 3 (C)) by simply creating a rectangular region around each pocket in which specular reflections were ignored.

After detecting the positions of each ball, ball colours must be determined. Again the fact that the environment of a snooker table is closed is an advantage here as there are only a small set of possible colours that can be present (white, red, yellow, green, brown, blue, pink and black). The average RGB colour in the environs of each detected ball is calculated and compared against a look-up table to determine the colour of each ball. The values in the look-up table were determined through experimentation.

Creating A Virtual Snooker Table

OpenGL (www.opengl.org) was used to create the virtual model of a snooker table used as part of this system. Figure 4 shows screenshots of an empty table and a table containing snooker balls which were created using simple OpenGL primitives. This is one aspect of the project which needs a considerable amount of further work as will be discussed in section 4.

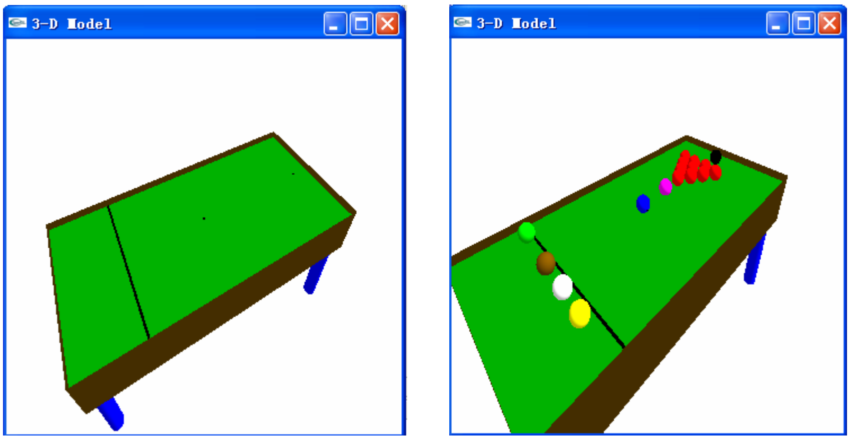

Fig. 4. The simple virtual snooker table used, shown empty and with balls.

4. Evaluation

In order to evaluate our system we performed a series of experiments in which various situations were set up on a small demonstration snooker table and the resultant virtual model produced by the system was compared against the real situation. This comparison sought to answer the following three questions:

- Did the system recognise all of the balls?
- Did the system recognise any spurious non-existent balls?
- Did the system get all ball colours correct?

In this evaluation 13 situations were used. Figure 5 shows a selection of these evaluation situations and the resultant 3D models. As can be seen, both situations likely to arise in a typical frame of snooker, and more exotic situations were used.

In the 13 evaluation situations used a total of 181 balls appeared. In all of these situations only one ball was not detected. The reason for this was that the missing ball was placed too close to the pocket area. The system was designed to filter out balls located close to the pockets as these were expected to be spurious reflections caused by the material inside the pockets. The solution to this is to use a more realistic shape around the pocket region rather than the rectangle used at present.

Over the course of the 13 situations evaluated two spurious extra balls were detected. The cause of these was dirt lying on the surface of the table. It is believed that by adding some more domain information to the detection system these mistakes could be avoided – or possibly the table could be kept clean!

Colour detection performed reasonably well, only making a mistake in 3 out of 181 detection tasks. The performance of the colour detection system could very easily be improved by reviewing the look-up table and making it more comprehensive.

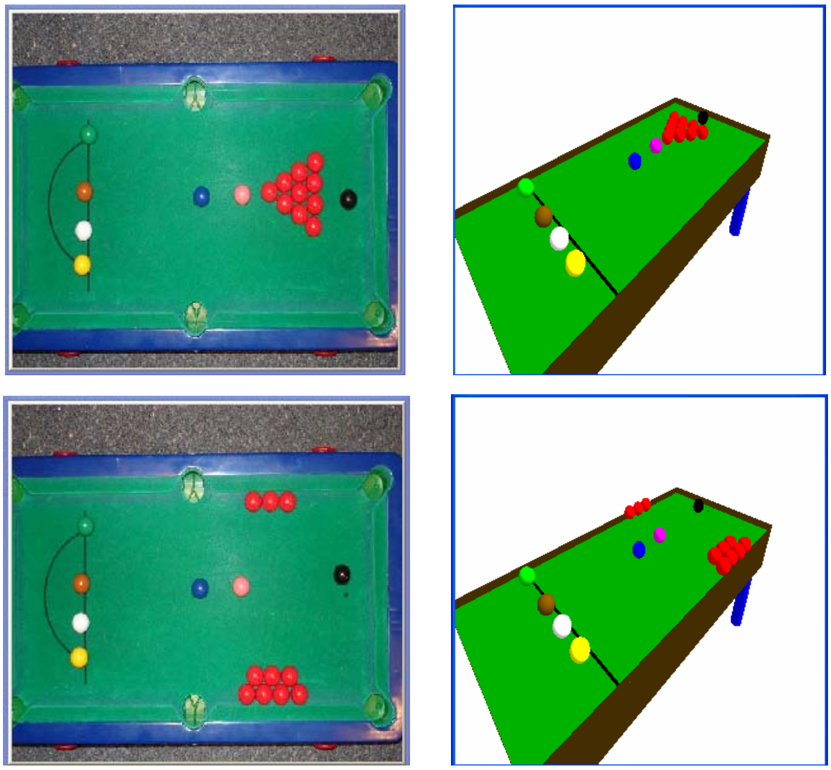

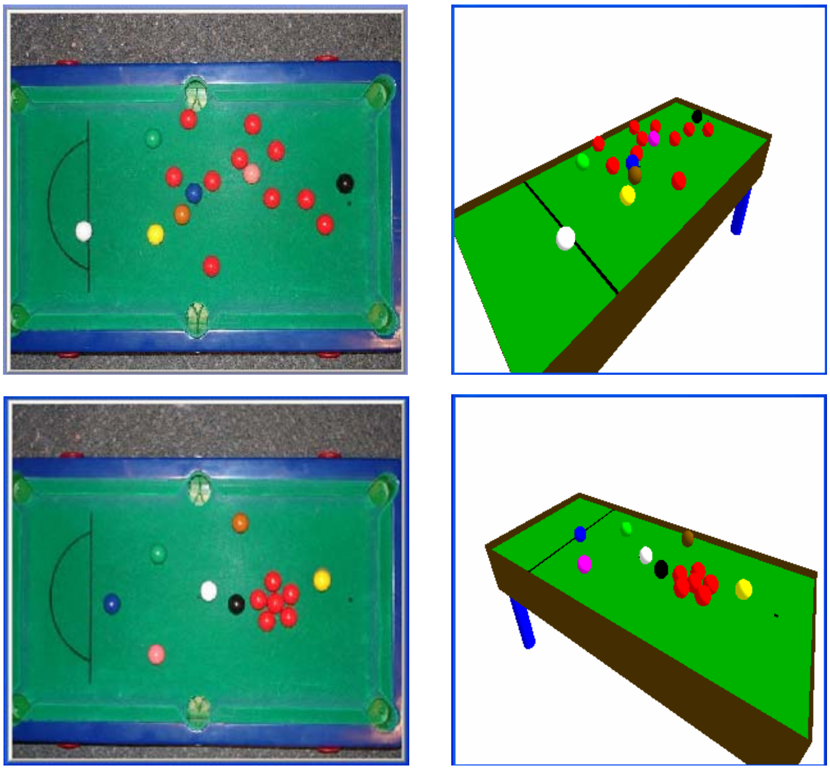

Fig. 5. Examples of images of a snooker table and the resulting virtual representation.

5. Conclusion

This paper has described the design, development and evaluation of the Snooker Extraction and 3D Builder (SE3DB) system, a system created to create a virtual 3D representation of a snooker table from a single overhead image of the table. The aim of the system is to render a 3-D model which can be used in televised snooker competition broadcasting to allow audiences view a rendering of the table from any angle. In developing the system, the key observation which made it possible was the realisation that each ball has a strong specular reflection that is very easily identified. The system has been evaluated and in most cases performs to an acceptable level of accuracy.

In the future we would hope to improve the system in the following ways:

- Create more realistic pocket regions so that balls placed next to pockets will not be filtered out as noise.

- Improve the colour detection system by improving the quality of the look-up table used to assign identified balls a colour.
- Create a much more visually appealing table model for the virtual 3D table renderings.
- Perform evaluations on a genuine competition snooker table.
- Add a calibration step to the system so that the system will cope with a camera positioned at any position and differences in specular reflection positions for balls on different parts of the table.

References

[1] Cavallaro, R. “The FoxTrax Hockey Puck Tracking System”, IEEE CG&A. vol. 17, no. 2, March/April, 1997

[2] Demiris, A. M.; Traka, M.; Reusens, E.; Walczak, K.; Garcia, C.; Klein, K.; Malerczyk, C.; Kerbiriou, P.; Bouville, C.; Boyle, E.; Ioannidis, N., “Enhanced sports broadcasting by means of augmented reality in MPEG-4”, In Proceedings of the International Conference on Augmented Virtual Environments and Three-Dimensional Imaging, 2001.

[3] Gonzalez, R.C.; Woods, R. E .“Digital image processing”, Prentice Hall, 2002.

[4] Khudeev, R., “A new flood-fill algorithm for closed contour”, In Proceedings of the IEEE International Siberian Conference on Control and Communications (SIBCON '05), pp172–176, 2005.

[5] Owens, N.; Harris C.; Stennett, C. “Hawk-eye tennis system”, In Proceedings of the International Conference on Visual Information Engineering: Ideas, Applications, Experience Visual Information Engineering (VIE2003), pp182-185, 2003.

[6] Pingali, G.; Opalach, A.; Jean, Y., “Ball Tracking and Virtual Replay for Innovative Tennis Broadcasts”, Pattern Recognition, In Proceedings. of the 15th International Conference on Pattern Recognition, Volume 4, pp152-156, 2000.

[7] Sharifi, M.; Fathy, M.; Mahmoudi, M.T.; “A classified and comparative study of edge detection algorithms”, In Proceedings of the International Conference on Information Technology: Coding and Computing, pp117–120, 2002.