Spatio-temporal Analysis of Tennis Matches

Silvia Vinyes Mora
Imperial College London
South Kensington Campus
London, UK
sv212@ic.ac.uk

William J. Knottenbelt
Imperial College London
South Kensington Campus
London, UK
wjk@doc.ic.ac.uk

ABSTRACT

Big data and analytics are having a major impact in sports – they have become essential for elite sports coaching, prediction and the enhancement of broadcasting. This has been possible because sophisticated technologies for collecting spatio-temporal data in sports have been developed. There is a real interest for sports analytics at the amateur level but unfortunately the availability of such technologies is still limited to the top tournaments due to their cost and requirements. In previous work we presented a low-cost technology for collecting tennis spatio-temporal data in real-time while being highly portable and non-intrusive. In this paper, we show that our system can also lead to valuable analysis of the game. Players and ball movements have been analysed separately and a combination of both has led to the extraction of higher level information. In contrast to previous work, we present an-end-to-end low-cost system for tennis analytics that goes from data collection to data analysis visualization.

Keywords

Sports data analysis; Tennis; Spatio-temporal analysis

1. INTRODUCTION

Sports analytics has always been of great interest but in recent years there has been a breakthrough on account of the rapid increase in the amount and detail of data that is available. The spatio-temporal data collected from sports events has been the major contributor to this data explosion. While for many years only statistical data could be analysed (eg. match scores, service games won, number aces etc.) now spatio-temporal data (eg. serve ball speed, distance run by each player) and high-level information (eg. net approaches) is also taken into account leading to a finer-grained analysis of the game.

The challenge addressed by this paper is how to analyse spatio-temporal data that has been generated by a system composed of commodity hardware.

KDD '16 August 14, 2016, San Francisco, CA, USA

DOI: 10.475/123_4

Although it is still an active area of research, sophisticated technologies exist for collecting spatio-temporal data in sports. Hawk-Eye is a state of the art technology for collecting highly accurate spatio-temporal data in a number of sports. In tennis, Hawk-Eye has even officially become an integral part of the game by assisting line-calls. However, the complexity and cost of this system make it only available to the main events at major tournaments. Therefore the data obtained with such systems is rich but limited in the number of matches and of difficult access. For instance the only published analysis from Hawk-Eye tennis data is of a single tournament [6, 7].

With a system like Hawk-Eye, one stores 8 values per frame (3D coordinates for each player and the ball); with a frame rate of 1 000fps this corresponds to 8 000 values per second. Extracting valuable information from such a large dataset is a real challenge. The objective is to present the data in a simple, meaningful and appealing way without disregarding important information for the analysis [1, 4, 3]. In [6, 7] the authors shown that spatio-temporal data obtained from Hawk-Eye, even though of a small sample, was very powerful in understanding patterns of the game and gaining insights into future events. However, the analytics that can be derived with such rich data is yet to be further explored.

We earlier presented a non-obtrusive system able to collect tennis spatio-temporal data by using commodity hardware and that can be installed on court in minutes [5]. Similarly to Hawk-Eye, the data that we collect are the positions of the centre of gravity of the players and that of the ball. In this paper we show that data collected from an accessible technology can also lead to an interesting analysis of the game; we also focus in the visualization aspect of the analysis.

2. METHODS

2.1 Data collection

Hardware and software

The system used to collect the data is composed of highly-portable commodity hardware. It comprises four high frame rate (60 fps), high definition, grey scale cameras and a GPU-enabled desktop computer.

The Pylon Camera Software was used to interface the cameras with the computer. A multi-threaded image processing application was implemented in C++ and OpenCV [2] to detect the players and the ball from the four cameras simultaneously and in real time. The frame-by-frame player and ball position is stored in a different file for each camera. Finally, the 2D coordinates extracted from the four cameras are triangulated to obtain the 3D position of each element(i.e. the ball and each player). This data is also stored in a file for further analysis and is used to represent the players and the ball in an interactive 3D virtual environment. A detailed evaluation of the system can be found in [5].

Data analysed

About 240 videos were analysed, corresponding to 60 different scenarios since each scene is captured by 4 cameras simultaneously and the videos range from 300 to 1500 frames each. It is important to note that the 2D and 3D data obtained from these videos contains some noise due to inaccuracies in the detection or non-detected elements.

Data cleansing

Each element in the ball data store contains the 3D coordinates of the detected ball position and its associated frame number. From the spatial distance between two consecutive elements (using 3D coordinates) and their time difference (from the frame number) outliers are detected and removed. These correspond to objects moving unrealistically fast, that is above a threshold based on the maximal velocity that can be reached by the tennis ball.

2.2 Data analysis

Heatmaps

Heatmaps are used to visualize the frequency of a player or ball being in a particular area of the tennis court, using 2D data (excluding the height dimension).The court and its immediate surrounding area are divided into a grid of equally sized rectangles (100×100 rectangles for the players and 200×200 for the ball) and then time spent by the player or ball in that area is shown by a color representing the normalized frequency.

Distance

The distance travelled by the players is calculated as the 2D Euclidean distance where the distance between two points a and b is:

$$\sqrt{(a_x - b_x)^2 + (a_y - b_y)^2}$$

In the case of distances between the player and the ball, distances are calculated in 3D and the distance between two points a and b is:

$$\sqrt{(a_x - b_x)^2 + (a_y - b_y)^2 + (a_z - b_z)^2}$$

Area covered

To represent the area covered by a player we divide the court and and its immediate surrounding area into a 100×100 grid and count the number of rectangles in the grid where the player has been. This is shown as a percentage of the total area.

Cubic fit

To detect when the ball is hit by a player we calculate the distance between the ball and player. The resulting data has missing values, mostly due to the removed outliers, and a cubic interpolation is performed to recover the missing data. Cubic fitting was chosen because tennis balls are known to decelerate quickly due to their surface coating.

3. RESULTS

3.1 Player movement profiles

Figure 1 is a heatmap of the players positions across all the videos analysed. This shows that players spend most of their time in the center behind the baseline, occasionally coming closer to the net. This pattern is in accordance to what we observed in the videos and what could be expected as the most common areas where a tennis players stands in a training scenario. This heatmap can also be generated for shorter spans of time to see trends in particular situations, to compare patterns across different players and to understand player interactions.

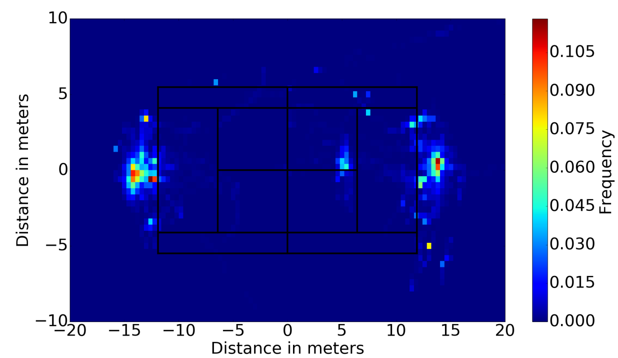

Figure 1: Players' position heatmap

3.2 Tennis ball

Figure 2 is a heatmap of the areas traversed by the tennis ball. It shows that the most common position of the ball is above the court near the center of the net. As seen in the videos, the ball is at different angles but crosses the net close to the middle and diverges to a less localised area when getting further away from the net.

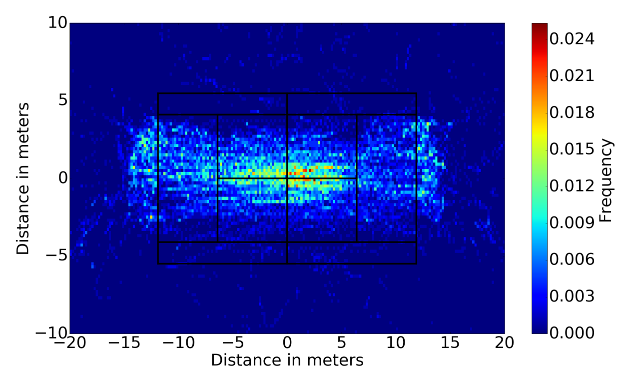

Figure 2: Stroke pattern density

3.3 Action: players and tennis ball combined

In this section, the data from the players is combined to that of the tennis ball in order to take the analysis a step

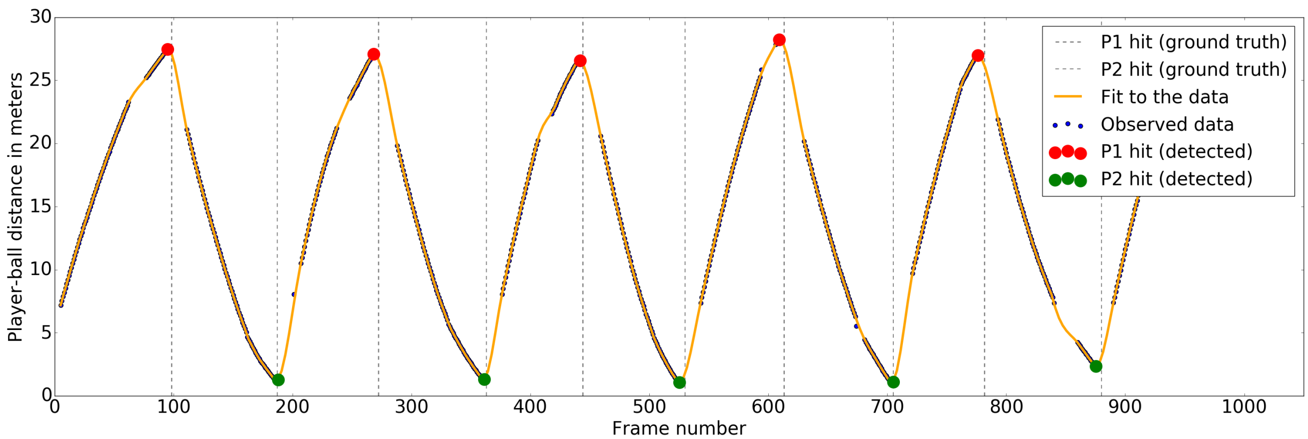

Figure 3: Distance between the player and tennis ball

Stroke number	1^{st}	2^{nd}	3^{rd}	4^{th}	5^{th}
Measured timing (in frames)	95	268	442	609	776
Real timing (in frames)	98	272	444	613	781
Error (in frames)	3	4	2	5	5
Error (in ms)	50	66	33	83	83

Table 1: Player 1 stroke detection

Stroke number	1^{st}	2^{nd}	3^{rd}	4^{th}	5^{th}
Measured timing (in frames)	188	361	525	705	875
Real timing (in frames)	187	363	530	705	880
Error (in frames)	1	2	5	0	5
Error (in ms)	16	33	83	0	83

Table 2: Player 2 stroke detection

further; here we present the results for one of the videos but other videos gave similar output. Figure 3 is a scatter plot of the player–ball 3D distance for one of the players using the automatic video detections.

Values close to zero in the y axis correspond to the ball being close to the current player and they can be interpreted as the frames in which he is hitting the ball. Similarly, maximal y values can be seen as the ball being hit by the other player. In order to find these local maxima and minima in data with missing values we performed a cubic interpolation, plotted in orange in Figure 3, and then we found the local maxima and minima from the interpolated data, shown in Figure 3 as red and green circles respectively. Their x coordinates correspond to the frame numbers at which strokes by one of the players are thought to occur.

To evaluate these results we manually collected the frame numbers at which the ball is hit by each of the players; this is represented by the vertical lines in Figure 3. Tables 1 and 2 show the comparison between the ground truth data and our automatic action detection. We can see that in the current example, the error in the timing at which an action occurs is between 0 and 5 frames corresponding to a maximal error of 83ms in real time, with a mean error of 53ms.

3.4 General statistics

Finally, we have captured a number of general statistics to enable the comparison between players in terms of:

- The average player speed when moving on the tennis court
- The percentage of court surface covered by each player
- The average speed of the strokes initiated by each player

Table 3 presents this information for 10 of the videos analysed.

These statistics are very valuable for understanding a player's trends and behaviour and differences amongst different players. In video number 77 for example, even though both players have a similar average speed, player 2 covers twice as much court area as player 1. This is illustrated in Figure 4, which shows the trajectories effectuated by each player in that video and where we can see that player 2 covered a larger court area by entering the service box.

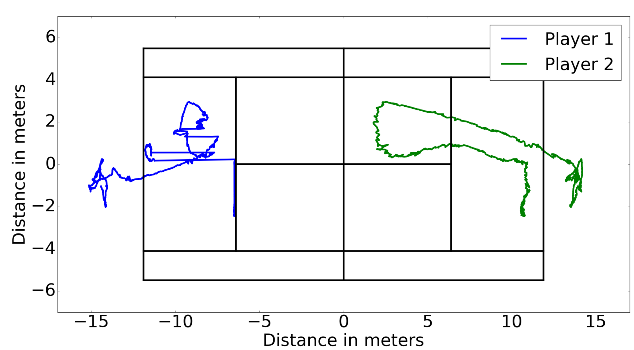

Figure 4: Players' trajectories

Video number	Average speed		Area coverage		Stroke average speed		Video length
	P1 (in meters/s)	P2 (in meters/s)	P1 (in %)	P2 (in %)	P1 (in meters/s)	P2 (in meters/s)	(in s)
25	0.8	1.7	1.1	1.4	21.1	17.4	14.2
27	1.1	1.7	1.2	1.6	30.5	29.2	8.2
36	11.5	1.2	6.0	1.1	16.8	33.0	14.2
40	24.2	0.9	6.6	0.8	31.2	29.8	11.2
44	0.9	1.0	1.6	1.5	21.6	22.7	13.7
53	0.7	0.4	0.4	0.2	30.3	40.9	6.6
64	1.1	3.0	2.7	7.4	39.5	31.7	21.7
68	0.8	1.3	1.8	2.2	25.6	30.0	19.4
70	2.1	0.9	5.1	2.4	29.3	31.3	23.5
77	1.7	1.9	5.8	10.8	28.2	25.3	28.4

Table 3: General statistics

4. CONCLUSIONS

This work, in conjunction with our system presented in [5], presents an end-to-end pipeline for the spatio-temporal analysis of a tennis match using a low-cost data collection system. The technology proposed is able to:

1. Collect spatio-temporal information in real-time using non-obtrusive, low-cost and highly portable hardware
2. Replay the events in a virtual 3D environment
3. Offer an appealing visualisation of the analysis of the data

The specific contribution of this paper is to show that the data that has been collected with a low-cost system can be used to represent patterns of player movement and ball trajectories and to extract high level information such as the players' trajectories or the timing of the strokes of each of the players.

5. FUTURE WORK

We have two main future work directions. First, add contextual information to the analysis to derive player and event specific information, as for example comparing the serve of two players. Second, look into specific interactions between players in addition to the aggregated data analysis shown in this article. In addition to this, we are planning to evaluation the stroke detection for more videos.

6. ACKNOWLEDGMENTS

This project has been funded by the EPSRC. I would like to thank G. Barnett, Casper da Costa-Luis, J. Garcia Maegli, M. Ingram and J. Neale for their contribution in developing the data collection technology and to the reviewers for their insightful comments.

7. REFERENCES

[1] A. Bialkowski, P. Lucey, P. Carr, Y. Yue, S. Sridharan, and I. Matthews. Large-scale analysis of soccer matches using spatiotemporal tracking data. In *Data Mining (ICDM), 2014 IEEE International Conference on*, pages 725–730. IEEE, 2014.

[2] G. Bradski. *Dr. Dobb's Journal of Software Tools.*

[3] D. Demaj. Geovisualizing spatio-temporal patterns in tennis: An alternative approach to post-match analysis. In *Proceedings of the 26th International Cartographic Conference*, 2013.

[4] G. Pingali, A. Opalach, Y. Jean, and I. Carlbom. Visualization of sports using motion trajectories: providing insights into performance, style, and strategy. In *Proceedings of the conference on Visualization'01*, pages 75–82. IEEE Computer Society, 2001.

[5] S. Vinyes Mora, G. Barnett, C. O. da Costa-Luis, J. Garcia Maegli, M. Ingram, J. Neale, and W. J. Knottenbelt. A low-cost spatiotemporal data collection system for tennis. In *Proceedings of the conference on Visualization'01*, pages 75–82. IEEE Computer Society, 2001.

[6] X. Wei, P. Lucey, S. Morgan, and S. Sridharan. Predicting shot locations in tennis using spatiotemporal data. In *Digital Image Computing: Techniques and Applications (DICTA), 2013 International Conference on*, pages 1–8. IEEE, 2013.

[7] X. Wei, P. Lucey, S. Morgan, and S. Sridharan. sweet-spot: Using spatiotemporal data to discover and predict shots in tennis. MIT Sloan Sports Analytics Conference, 2013.