CrossMark

ORIGINAL ARTICLE

Robust ultrasound probe tracking: initial clinical experiences during robot-assisted partial nephrectomy

Philip Pratt[1] · Alexander Jaeger[2] · Archie Hughes-Hallett[3] · Erik Mayer[3] ·
Justin Vale[3] · Ara Darzi[1,3] · Terry Peters[4] · Guang-Zhong Yang[1]

Abstract

Purpose In order to assist in the identification of renal vasculature and tumour boundaries in robot-assisted partial nephrectomy, robust ultrasound probe calibration and tracking methods are introduced. Contemporaneous image guidance during these crucial stages of the procedure should ultimately lead to improved safety and quality of outcome for the patient, through reduced positive margin rates, segmental clamping, shorter ischaemic times and nephron-sparing resection.

Methods Small KeyDot markers with circular dot patterns are attached to a miniature pickup ultrasound probe. Generic probe calibration is superseded by a more robust scheme based on a sequence of physical transducer measurements. Motion prediction combined with a reduced region-of-interest in the endoscopic video feed facilitates real-time tracking and registration performance at full HD resolutions.

Results Quantitative analysis confirms that circular dot patterns result in an improved translational and rotational working envelope, in comparison with the previous chessboard pattern implementation. Furthermore, increased robustness is observed with respect to prevailing illumination levels and out-of-focus images due to relatively small endoscopic depths of field.

Conclusion Circular dot patterns should be employed in this context as they result in improved performance and robustness. This facilitates clinical usage and interpretation of the combined video and ultrasound overlay. The efficacy of the overall system is demonstrated in the first human clinical case.

Keywords Ultrasound · Tracking · Image guidance · Partial nephrectomy · Robot-assisted

✉ Philip Pratt
p.pratt@imperial.ac.uk

1 Hamlyn Centre for Robotic Surgery, Imperial College, London, UK

2 Electrical and Electronic Engineering, University College, Cork, Ireland

3 Department of Surgery and Cancer, Imperial College, London, UK

4 Robarts Research Institute, University of Western Ontario, London, Canada

Purpose

The image-enhanced approach to surgical intervention aims to improve patient safety and oncological outcome, and this is demonstrated particularly well in robot-assisted partial nephrectomy [1], where an accurate appreciation of the typically variable renal vasculature and tumour interface is of paramount importance. The provision of an image guidance system based on preoperative imaging alone goes some way to facilitating this goal [2], in the sense that the topological relationships between structures are relatively straightforward to illustrate macroscopically. But to reach full efficacy, such a system must be complemented by a contemporaneous imaging modality [3], immune to the undesirable effects of tissue deformation and capable of resolving structure at a more detailed level. Ultimately, this should lead to reduced positive margin rates, the option of segmental clamping, shorter ischaemic times and nephron-sparing resection.

To this end, the use of intracorporeal 'pickup' ultrasound transducers [4–6] is becoming prevalent, where the resulting

2D scan images are automatically registered and presented in real time within the 3D surgical scene. Recent studies [4,7] have considered conventional tracking methods (e.g. electromagnetic or optical tracking of the distal probe), but the ubiquity of ferromagnetic materials and awkward lines of sight in the operating theatre present problems that invariably hinder translation of this technology into everyday clinical practice. The endoscope itself can double as a tracking device, thereby eliminating a number of intermediate coordinate systems and, therefore, sources of registration error. Visual identification of a geometric pattern rigidly attached to the probe is sufficient to estimate its pose, ultimately leading to direct registration [8,9] of the ultrasound image.

Even so, the problem of how to create an accurate ultrasound probe tracking algorithm of this nature, sufficiently fast and robust to work in a live and unforgiving clinical environment, has yet to be solved adequately. These criteria are prerequisites for widespread adoption of the technology. Therefore, learning from clinical experience, this study extends previous work [10] by adopting the principle of simplicity, and introducing a tracking pattern made up of circular dots. It also considers the output of the system, where the psychophysical effects of the method chosen to display the registered and fused images play a crucial role in anatomical interpretation and ease of use of the surgical guidance platform.

Methods

Circular dot tracking

The marker used for tracking was an asymmetric circle KeyDot® pattern (Key Surgical Inc., Eden Prairie, MN, USA). The marker was initially attached to a 'dummy' probe for the purpose of testing. Figure 1 illustrates the design of the pattern comprising an asymmetric grid of circles with three rows and seven columns. Vertical and horizontal dot spacing was measured using Vernier calipers at 1.19 mm, with each dot having a diameter of 500 μm. This pattern design was chosen as dots are geometrically simpler than squares, and
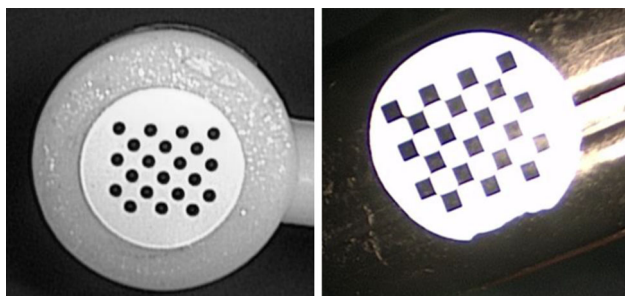
Fig. 1 Asymmetrical circular dot pattern (*left*) and original chessboard pattern (*right*)

it was anticipated that circular structures are less affected by degrading image effects such as blooming and smudging (Fig. 1, right). The tracking algorithm for detecting circles does not rely on the intersection of edges, but instead relies on an efficient 'blob detector' approach. Intuitively, tracking dots should be more robust when the reliance on well-defined edge crossings is eliminated.

To facilitate repeatability, the OpenCV implementation of the 'simple blob detector' was employed [11]. At each invocation of *findCirclesGrid()*, the following steps are executed sequentially:

- The source colour image is converted to a sequence of binary images by applying greyscale thresholds at equally spaced levels between specified minimum and maximum bounds.
- Connected components are extracted from each binary image using a contour extraction routine [12] utilising border following. The centres are found by computing the moments of each rasterised image.
- Blob centres are grouped by their coordinates across the sequence of thresholded images, in accordance with a specified 'minimum distance between blobs' parameter. Sufficiently large groups correspond to positively identified blobs.
- The final centre positions are estimated by a weighted mean. The use of an asymmetric pattern ensures that the sorted results are returned in a consistent manner.

The relevant detector parameters were set as follows: $minThreshold = 10$; $maxThreshold = 220$; $minDistBetweenBlobs = 5$ pixels; $filterByArea = $ TRUE; $minArea = 25$ pixels^2; $maxArea = 5000$ pixels^2; $filterByIntertia = $ TRUE; $minInertiaRatio = 0.1$; $filterByConvexity = $ TRUE; $minConvexity = 0.95$; $filterByColor = $ TRUE; $blobColor = 0$; and $minRepeatability = 2$.

Tracking regions of interest

The OpenCV *findCirclesGrid()* function is not accelerated by any means, in the sense that the internal components of the library function do not make use of parallel processing constructs nor heterogeneous computing APIs. This results in prohibitively long processing delays between frames when high-definition (HD) video streams are prevalent. A continuously updating cropping technique was developed to reduce the number of pixels that the function was required to process. On detecting the circular dot pattern initially, a set of rectangular coordinates whose area is proportional to the size of the pattern are calculated such that the coordinates envelop the location of the pattern in the image plane (Fig. 2). This set of coordinates is used to crop the next video frame prior to calling *findCirclesGrid()*. The cropping process is repeated for
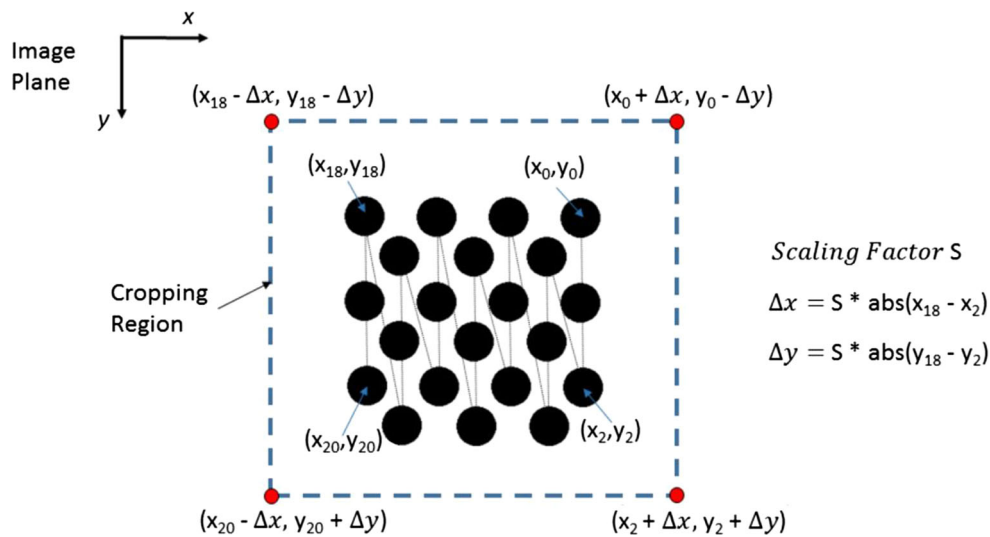
Fig. 2 Calculating the appropriate cropping area around the circular dot pattern

every frame. The method significantly reduces the number of pixels required for processing as it removes the majority of redundant information from the image. The relative size of the cropping area can be adjusted using the scaling factor S. A higher scaling factor yields a larger cropping area, at the expense of a degradation in performance. Chosen empirically, a scaling factor of $S = 0.4$ was used during the experiments herein.

However, generation of a cropping region for use in future video frames inevitably causes failures when the pattern moves sufficiently far away. Since the cropping region calculation did not initially take into account the current velocity of the pattern, resulting in partial patterns and failures in recovering the location of all circular dots in the pattern, an additional step was included to account for such motion. Figure 3 illustrates the salient features of the calculation. The velocity of the pattern is estimated by taking the vector difference of the positions of the two previous consecutive cropping rectangles. The resulting displacement vector is multiplied by a scaling factor S_v (set empirically to 0.5) and accumulated with the following cropping rectangle, thus amending the calculation with an approximate velocity-based 'prediction' of the location of the dot pattern in the next frame. Should the circular dot detection fail at any point, e.g. on account of a rapid motion, the region-of-interest is reset, a threshold step size of 30 is adopted, and the entire image is processed. Otherwise the default size 10 is employed, allowing for finer steps during blob detection, permissible in the knowledge that the image has been cropped.
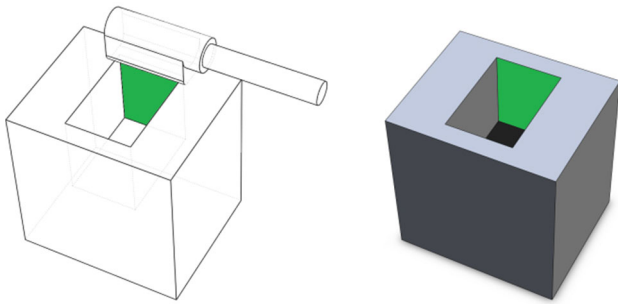
Ultrasound overlay

Once the probe pattern has been recovered, its pose relative to the camera coordinate system can be estimated using the
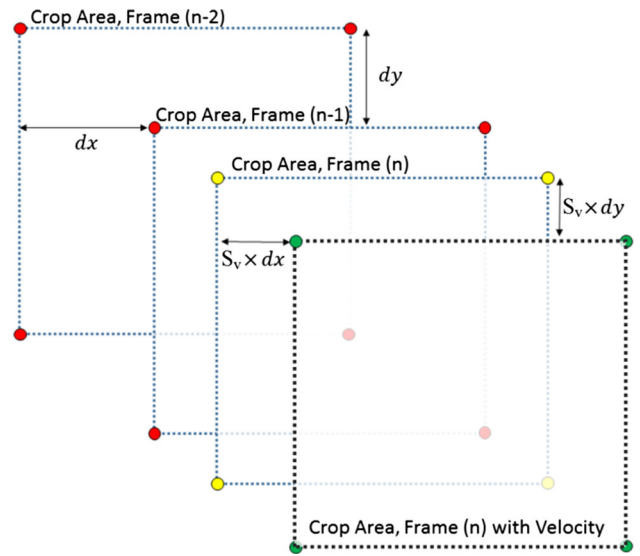
Fig. 3 Cropping rectangles with velocity estimation and weighted accumulation

camera's calibrated [13] intrinsic properties. This transformation is concatenated with the constant image-to-pattern relationship (i.e. probe calibration) in order to determine the registered position of the ultrasound 'slice' within the 3D surgical scene. However, within the stereo display, the precise means by which the left and right operative video feeds are combined with the ultrasound texture will have a significant impact on the surgeon's ability to appreciate any observed subsurface structure.

Three options were made available in this study. In each case, the ultrasound image is rendered into the left and right views of the surgical scene in accordance with the stereo camera intrinsic and extrinsic parameters. The first

Fig. 4 Geometry of windowed overlay option showing transducer and back wall (*green*)

two options comprise overwriting and alpha-blending of the ultrasound texture. The third, and perhaps most compelling, option requires that a window, perpendicular to the ultrasound image, is removed from the original left and right images. Figure 4 illustrates the underlying geometrical arrangement with wireframe and solid views of a simple model. The window is also rendered into the alpha channel of the output buffers, such that per-pixel blending operations can be used subsequently to clip the 'vertical' quadrilaterals. The 'back wall' of the ultrasound image is rendered in this manner, together with any measurement gradations. Finally, any visible 'side walls' of the hole are added—again, with measurement gradations if enabled. Together, these features and their respective realisations in the left and right camera views help give the observer the illusion that part of the target organ's surface and interior have been cut away [14]. Employing PVA cryogel kidney phantom models, Fig. 5 illustrates the overlay options monoscopically.

Probe calibration

Preliminary work [8] described a general probe calibration method applicable to the UST-533 microsurgery probe. In total, seven degrees of freedom were determined, ultimately resulting in the transformation that takes pixels in the ultrasound image to the correct position in the probe pattern coordinate frame. They comprised an image scale factor and a rigid 3D rotation and translation.

While in practice the method was capable of yielding very accurate calibrations, the relatively high dimensionality of the parameter space meant that the results were very sensitive to small deviations in the input data, e.g. the centres of the Z-phantom wire images in the reference scans. A more robust scheme was required that exploited features of the ultrasound cart and the relatively simple geometry of the transducer casing in this instance.

To this end, the pixel-to-millimetre scale factor is now determined directly using the ultrasound cart's built in measurement capability. Careful alignment of the circular dot pattern axes with respect to the rectilinear probe casing ensures that the rotation component of the calibration can be set to the identity. The X and Z components of the translation are determined by direct measurement of the transducer position relative to its outer casing. The remaining degree of freedom, the Y translation, is determined by alignment of the direct view of a single cross-wire phantom feature [8] with its ultrasound image overlay.

Probe decontamination and deployment

Driven by a ProSound α-10 cart, the ultrasound transducer used in this study was a UST-533 multifrequency linear array microsurgery probe (Hitachi Aloka Medical Ltd., Tokyo, Japan). It is certified for sterilisation using a hydrogen peroxide plasma process, and in this instance, a STERRAD 100NX unit was employed (Johnson & Johnson, New Brunswick, NJ, USA). In order that the transducer can be picked up by da Vinci *ProGrasp*^TM or Cadiere forceps instruments (Intuitive Surgical, Sunnyvale, CA, USA), a small friction fit probe clip was designed and 3D-printed in implant-grade cobalt–chrome–molybdenum superalloy using the direct metal laser sintering (DMLS) process (Fig. 6). This material is suitable for high-pressure steam sterilisation in an autoclave.

To validate the sterilisation processes, the ultrasound probe (with KeyDot markers attached to both faces), probe

Fig. 5 Ultrasound overlay schemes: overwrite (*left*), alpha-blend (*middle*) and window (*right*)
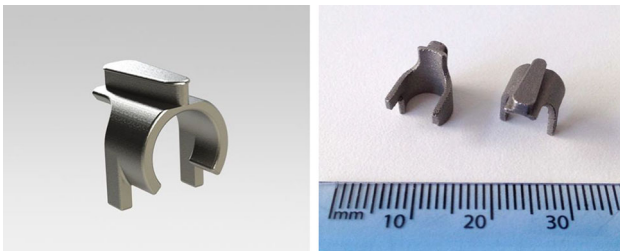
Fig. 6 Ultrasound probe clip design (*left*) and manufactured instances with scale (*right*)

Fig. 7 Tether suture being applied to ultrasound probe clip prior to insertion

clips and holding trays were deliberately contaminated and then subjected to their documented washing and sterilisation cycles. Swab samples were taken at the reprocessing facility, both before and after cleaning, and were sent for independent microbiological analysis at an accredited laboratory, in accordance with the ISO 11137 standard. For each item and the control, the number of colony forming units (CFUs) was measured to be well within expected limits.

Acting as an additional safety measure, a long suture is attached to the ultrasound probe clip before insertion into the patient's body through the assistant port (Fig. 7). In the unlikely event that the clip becomes detached from the transducer, the tether can be used to retrieve it through the port.

Results

Probe calibration

To assess probe calibration accuracy, eight different standard-definition camera and transducer poses were used to compare ultrasound images of the validation phantom crossing point [8] against the corresponding direct line-of-sight views, afforded when the phantom is drained of water with its door removed. The mean, standard deviation and maximum of the overlay errors were found to be 537, 90 and 653 μm, respectively, and can be seen to be an improvement over

Table 1 Operational envelope for circular dot marker pattern

	Range (mean)	Min (mean)	Max (mean)	Max (SD)
Rotate X	102.5°	−50.1°	52.4°	4.1°
Rotate Y	110.3°	−52.7°	57.6°	3.4°
Translate Z	73 mm	13 mm	86 mm	4.6 mm

the validation results (669, 255 and 961 μm, respectively) generated by the previous 7-DOF optimisation method. For comparison, a mean stereo camera calibration reprojection error of 0.50 pixels was recorded.

Operational envelope

Physical measurements were taken to determine the operational envelope of the ultrasound probe to which the circular dot pattern marker had been attached. The results are summarised in Table 1. The minimum and maximum attainable rotations and translations in the X, Y and Z planes (each with respect to the camera's coordinate system) were recorded. Each minimum and maximum measurement was taken ten times. Rotational measurements about the X and Y axes were recorded at a distance of 55 mm from the endoscope tip. A camera exposure time of 1/85th of a second was used. The light fountain (Karl Storz GmbH, Tuttlingen, Germany) was set at an intensity level of 75 % for the duration of the experiment. A Wolf stereo endoscope (Richard Wolf GmbH, Knittlingen, Germany) with HD camera heads was utilised in the laboratory setting.

Through comparison with the operational envelope results generated for the original chessboard pattern implementation [8], the extent of the incremental improvement can be calculated. An increase of 5.7 and 8.1 % in the X- and Y-axis rotation ranges is observed, respectively. Whereas these changes are relatively modest, a substantially large increase of 74 % in the Z translation range is also observed. This is perhaps the most significant measure of the operational envelope, as the surgeon will naturally try and orient the ultrasound probe such that the tracking pattern and hence the registered ultrasound image itself are perpendicular to the long axis of the endoscope.

Furthermore, qualitative observations were made by using the circular dots and chessboard pattern markers in a side-by-side configuration, such that footage could be recorded and then played back through the respective tracking algorithms. The use of the chessboard marker required more careful placement from the user with regard to its position in front of the camera. A more robust experience was found using the circular dot marker as its behaviour was less sensitive to large position changes.

Illumination levels

Endoscopic videos of the chessboard and circular dot patterns in the same physical location were recorded, comprising a sequence of ten increasing illumination levels, equally spaced throughout the extent of the light fountain dial (Fig. 8), at three different distances from the endoscope tip: 15 mm (Near, as illustrated); 26 mm (Middle); and 38 mm (Far).

Figure 9 illustrates that at each representative distance, the circular dot pattern performance is superior. Sensor blooming has a more deleterious effect on the sharp corners of the chessboard, in the sense that they become disconnected as pixels bleed into each other. This is in contrast to a circle, which will remain topologically identical until it disappears completely. Storing the last best tracking transformation can help overcome intermittent tracking behaviour, but only up to a point. With respect to illumination, the circular dot pattern behaves more robustly.
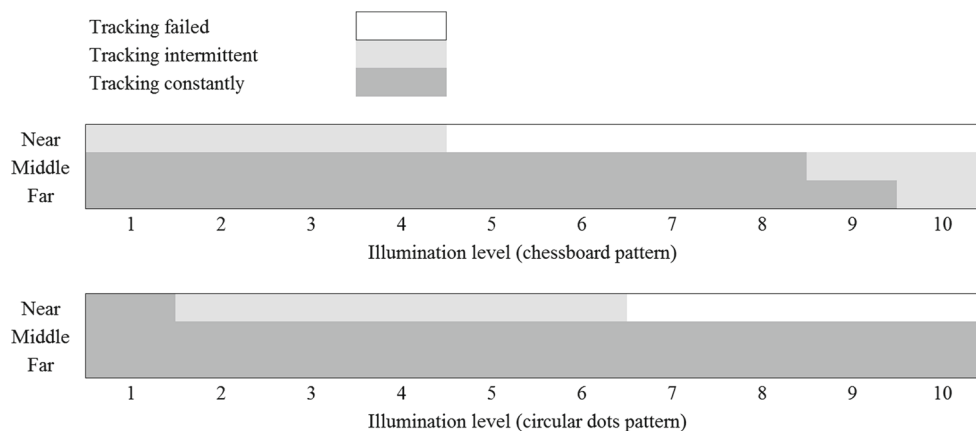
Focus position

The performance of the chessboard and circular dot pattern tracking algorithms with respect to changes in focus posi-tion (equivalently, with respect to the tracking target moving in and out of the endoscope's relatively shallow depth of field) bears a similar relation. It can be seen from Fig. 10 that the circular dot pattern has a wider successful tracking enve-lope. Closer inspection reveals that the latter algorithm is less prone to intermittent behaviour—it is therefore less suscep-tible to imperfections in the video input, which is inevitably noisy. Again, this can be attributed to the geometrically sim-pler nature of the pattern—there are no sharp corners to erode.

Typically, during the course of an intervention, there will be periods when even a high level of robustness with respect to the target focus quality will be insufficient to accommodate the necessarily changing spatial relationship between camera and surgical scene. It is therefore imperative that any practical system has both robustness and a mechanism for rapid recal-ibration when the extent of the scene change is too great. It is worth noting that the relationship between the ultra-sound image and probe pattern coordinate system remains constant, and therefore it is necessary only to repeat the endo-scope intrinsic parameter calibration. This can be achieved using the tracking pattern itself, in conjunction with a pre-calibrated model of endoscope focus characteristics across the entire focus range [15].
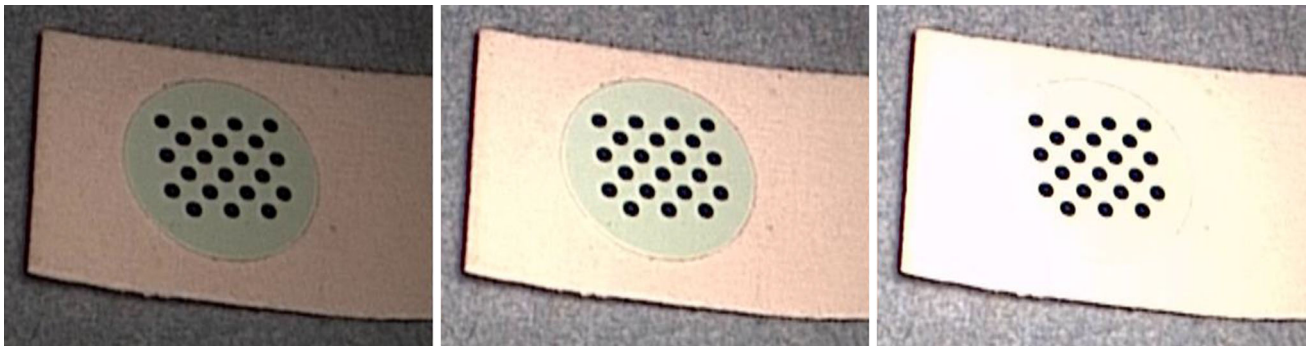
Fig. 8 Range of illumination levels (reaching saturation and blooming in the right-most image)

Fig. 9 Tracking results over illumination range (intermittent and continuous)
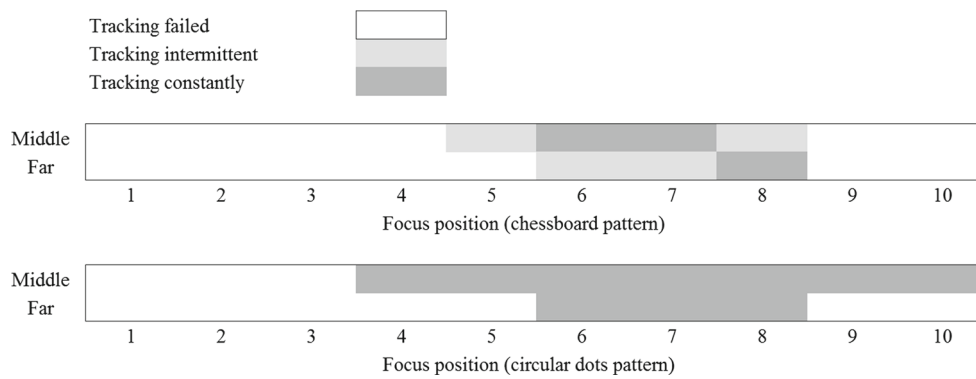
Fig. 10 Tracking results over focus position range (intermittent and continuous)

Table 2 Tracking performance comparison at different resolutions and distances

Resolution	Pattern	Far (ms)	Near (ms)	Empty (ms)	Failure (ms)
PAL	Circular dots	5	12	13	14
720 × 576	Chessboard	29	29	25	30
HD	Circular dots	11	22	75	60
1920 × 1080	Chessboard	17	21	11	22

Tracking performance

Minimal video processing delay is critical to the successful uptake and deployment of any image guidance system in clinical practice. Using different camera systems with different resolutions, and measuring tracking processing time on the live system (where the implementation is instrumented with high-resolution timers), Table 2 compares the performance of the original chessboard pattern versus the new circular dot pattern. In the low-resolution case (PAL), typical of many laboratory configurations, the blob detector algorithm with predictive cropping significantly outperforms the original chessboard approach. It should be noted that in this instance, unlike the HD case, the initial corner detector was not accelerated with CUDA [8] since without it, acceptable frame rates and delays are achievable at the lower resolution.

However, most systems currently in clinical practice will be using HD resolution, and so the more important set of results are to be found in the lower two rows. While the performance of the two algorithms is very similar when the tracking target is at the position closest to the endoscope, the predictive cropping comes into play significantly when the target is positioned at a distance (processing times highlighted in yellow—11 versus 17 ms). This is due to the fact that the size of the cropping window is proportional to the observed spacing between the circular dots. The further away the pattern, the smaller the image area that must undergo thresholding, contour delineation and blob detection.

The relatively fast processing time of 11 ms in the HD/chessboard case, where the image was empty (i.e. devoid of any pattern), can be attributed to the fact that the prevailing algorithm has an early exit clause which is invoked when the number of detected corners falls below a predefined threshold. Moreover, the high processing times of 75 and 60 ms, for the HD circular dots case, represent the delay that could be expected if predictive cropping were not enabled. Once tracking has been enabled, the processing will run at this speed until the algorithm 'locks on' to the dot pattern. This initial period is typically very brief and subjectively does not present any issues.
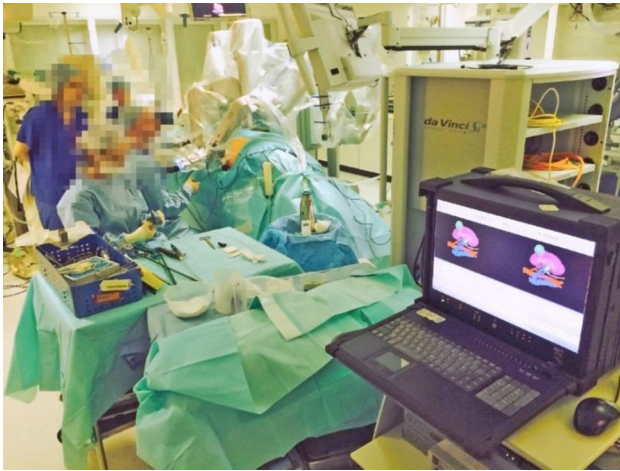
Human case report

The following report documents the first human use of circular dot patterns for tracking in ultrasound-guided robotic partial nephrectomy. The patient recruited to this study was a 67-year-old female with a tumour approximately 3.5 cm in diameter presenting in the left kidney. A preoperative CT scan of the urinary tract, with contrast, at high resolution, was acquired approximately 2 months prior to the date of the intervention, using a Brilliance 64-channel scanner (Philips Healthcare, Amsterdam, The Netherlands). The axial slice thickness was 2 mm, with an in-plane resolution of 512 by 512 voxels.

Informed and written consent was obtained from the patient prior to the intervention under the ethical protocol for the study, entitled 'Improving Outcomes in Robotic and Endoscopic Surgery using Augmented Reality Guidance'. The most recent substantial amendment to this study (REC reference 07/Q0703/24) was given a favourable ethical opinion by the NRES Committee London–Dulwich. The da Vinci stereo HD 30° endoscope was calibrated [13] at the outset of the procedure during the robot preparation phase. The mean reprojection error was recorded as 0.46 pixels.

The preoperative imaging was segmented using ITK-SNAP [16], and mesh-based renderings of the left kidney,
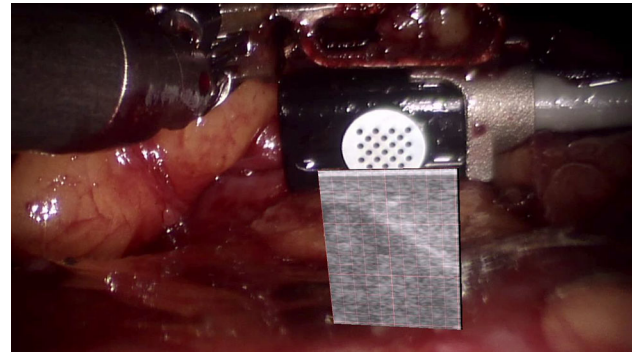
Fig. 11 Theatre layout showing the image guidance portable workstation in the foreground

Fig. 13 Still frame from left console feed illustrating registered image of tumour interface

tumour, vasculature and ureter were provided in the surgical console in anticipation of the intraoperative ultrasound probe deployment (Figs. 11, 12). The structure of the left renal hilar vasculature was relatively uncomplicated (single branch from the aorta, one proximal bifurcation and another distal bifurcation), which permitted the application of a single arterial clamp proximal to the aorta in order to achieve a period of warm ischaemia, during which the resection took place.

Although not directly pertinent to this study, these details are included in the case report to put the developmental process in a clear clinical context and to emphasise the staged approach where different imaging modalities are employed at different stages of the intervention. An integrated navigation platform such as this, offering multimodal, multiscale guidance can give the surgeon the most appropriate information (spatial, topological or otherwise) at the key junctures, to ensure the best patient outcomes.

Immediately prior to resection, the registered ultrasound overlay was used to confirm the location of the tumour. As can be seen in Fig. 13, the registered scan clearly demonstrates the interface between the normal parenchyma and tumour. The light red gradations on the image denote intervals of 1 mm. Anecdotal surgeon feedback immediately following the case suggested that while the accurately registered overlay capability was efficacious, the partial occlusion or blending of the surgical scene was a source of distraction from time to time. Negative tumour margins were reported postoperatively.

Conclusion

Both quantitative and qualitative results indicate that circular dot patterns are a significant improvement over the previous chessboard implementation, and a successful human trial heralds more extensive clinical evaluation. The observed level of robustness aids the development of 3D volume reconstructions based on freehand slice acquisition, which may help overcome some of the psychophysical barriers to interpretation. Paradoxically, the action of ultrasound scanning and overlay, while improving the surgeon's ability to appreciate the tumour location, impairs their ability to perform resection at the same time. Therefore, important research remains to be undertaken in this area.
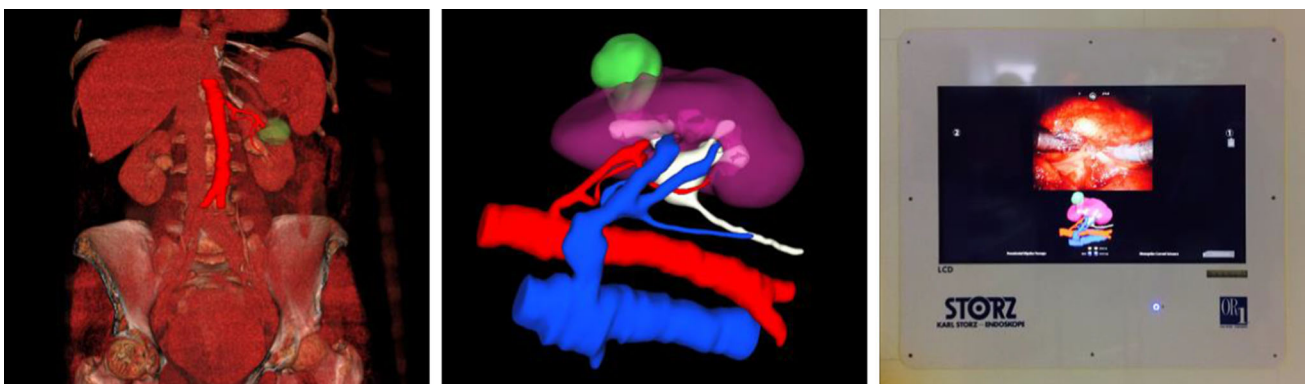
Fig. 12 Console and assistant views of left kidney showing vasculature and tumour location

One approach that has received some attention involves placing the ultrasound transducer in a location behind or alongside the organ of interest, where it is also invisible or only partially visible through the endoscope, but at a distance from the dissection site. This requires an unavoidable intraoperative registration procedure, for example, using a tracked tool to generate feature correspondences at the air-tissue boundary [17], or using a similar method employing the photoacoustic effect [18] where sound waves are generated following light absorption at visible points. Thus far, no solution represents a panacea; in this instance, the registration is dependent on the prevailing camera position.

While the core of the tracking algorithm described herein uses an open-source implementation, thereby promoting repeatability, calibration and tracking methods need to be optimised for the range of dedicated and commercially available robotic ultrasound probes now on the market. Furthermore, while the robot kinematics alone are insufficiently accurate to register the endoscopic and ultrasound views over an extended time period, a complementary scheme could mitigate those situations where tracking failures inevitably arise.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki Declaration and its later amendments or comparable ethical standards. This article does not contain any studies with animals performed by any of the authors.

Informed consent Informed consent was obtained from all individual participants included in the study.

References

1. Hughes-Hallett A, Mayer E, Marcus H, Cundy T, Pratt P, Darzi A, Vale J (2013) Augmented reality partial nephrectomy: examining the current status and future perspectives. Urology 83(2). Elsevier, 266–273

2. Hughes-Hallett A, Pratt P, Mayer E, Martin S, Darzi A, Vale J (2014) Image guidance for all—TilePro display of 3-dimensionally reconstructed images in robotic partial nephrectomy. Urology 84(1). Elsevier, 237–243

3. Sun M, Wagner A, San Francisco I, Brook A, Kavoussi L, Russo P, Steele G, Viterbo R, Pedrosa I (2012) Need for intraoperative ultrasound and surgical recommendation for partial nephrectomy: correlation with tumor imaging features and urologist practice patterns. Ultrasound Q 28(1):21–27

4. Leven J, Burschka D, Kumar R, Zhang G, Blumenkranz S, Xiangtian D, Awad M, Hager G, Marohn M, Choti M, Hasser C, Taylor R (2005) DaVinci canvas: a telerobotic surgical system with integrated, robot-assisted, laparoscopic ultrasound capability. In: MICCAI, LNCS 3749. Springer, Berlin, pp 811–818

5. Schneider C, Dachs G, Hasser C, Choti M, DiMaio S, Taylor R (2010) Robot-assisted laparoscopic ultrasound. In: Information Processing and in Computer-Assisted Interventions, LNCS 6135. Springer, Berlin, pp 67–80

6. Schneider C, Guerrero J, Nguan CY, Rohling R, Salcudean SE (2011) Intra-operative "Pick-Up" ultrasound for robot assisted surgery with vessel extraction and registration: a feasibility study. In: International conference on information processing in computer assisted interventions. Springer, Berlin, pp 122–132

7. Cheung CL, Wedlake C, Moore J, Pautler SE, Peters TM (2010) Fused video and ultrasound images for minimally invasive partial nephrectomy: a phantom study. In: MICCAI, Part III, LNCS 6363. Springer, Berlin, pp 408–415

8. Pratt P, Di Marco A, Payne C, Darzi A, Yang G-Z (2012) Intraoperative ultrasound guidance for transanal endoscopic microsurgery. In: MICCAI, Part I, LNCS 7510. Springer, Berlin, pp 463–470

9. Jayarathne U, McLeod AJ, Peters T, Chen E (2013) Robust intraoperative US probe tracking using a monocular endoscopic camera. In: MICCAI, Part III, LNCS 8151. Springer, Berlin, pp 363–370

10. Hughes-Hallett A, Pratt P, Mayer E, Di Marco A, Yang G-Z, Vale J, Darzi A (2013) Intraoperative ultrasound overlay in robot-assisted partial nephrectomy: first clinical experience. Eur Urol 65(3). Elsevier, 671–972

11. Bradski G, Kaehler A (2008) Learning OpenCV: computer vision with the OpenCV library. O'Reilly Media, Inc., Sebastopol

12. Suzuki S, Abe K (1985) Topological structural analysis of digitized binary images by border following. Comput Vis Graph Image Process 30(1):32–46

13. Zhang Z (2000) A flexible new technique for camera calibration. IEEE Trans Pattern Anal Mach Intell 22(11):1330–1334

14. Fuchs H, State A, Pisano E, Garrett W, Hirota G, Livingston M, Whitton M, Pizer S (1996) Towards performing ultrasound-guided needle biopsies from within a head-mounted display. In: Visualization in Biomedical Computing, LNCS 1131. Springer, pp 591–600

15. Pratt P, Bergeles C, Darzi A, Yang G-Z (2014) Practical intraoperative stereo camera calibration. In: Proceedings of the international conference on medical image computing and computer-assisted intervention. Part II, LNCS 8674. Springer, Berlin, pp 667–675

16. Yushkevich P, Piven J, Cody-Hazlett H, Gimpel-Smith R, Ho S, Gee J, Gerig G (2006) User-guided 3D active contour segmentation of anatomical structures. Neuroimage 31(3):1116–1128

17. Yip M, Adebar T, Rohling R, Salcudean S, Nguan C (2010) 3D Ultrasound to stereoscopic camera registration through an air-tissue boundary. In: MICCAI, LNCS 6362. Springer, Berlin, pp 626–634

18. Cheng A, Kang J, Taylor R, Boctor E (2012) Direct 3D Ultrasound to video registration using photoacoustic effect. In: MICCAI, LNCS 7511. Springer, Berlin, pp 552–559