Automated Resection Planning for Bone Tumor Surgery

Dave Hill [1], Tom Williamson [1], Chow Yin Lai [2], Martin Leary [1], Milan Brandt [1], Peter Choong [3]

[1] Centre for Additive Manufacturing, School of Engineering, RMIT University, 58 Cardigan St, Carlton, Australia, 3001

[2] Department of Electronic and Electrical Engineering, University College London, Malet Place and Torrington Place, Roberts Building, Level 7, London, United Kingdom, WC1E 7JE

[3] Department of Surgery, University of Melbourne, Level 2, Clinical Sciences Building, 29 Regent Street, Fitzroy, Australia, 3065

Abstract— Planning for bone tumor resection surgery is a technically demanding and time-consuming task, reliant on manual positioning of cutting planes (CPs). This work describes an automated approach for generating bone tumor resection plans, where the volume of healthy bone collaterally resected with the tumor is minimized through optimized placement of CPs. Particle swarm optimization calculates the optimal position and orientation of the CPs by introducing a single new CP to an existing resection, then optimizing all CPs to find the global minima. The bone bounded by all CPs is collaterally resected with the tumor. The approach was compared to manual resection plans from an experienced surgeon for 20 tumor cases. It was found that a greater number of CPs reduce the collaterally resected healthy bone, with diminishing returns on this improvement after five CPs. The algorithm-generated resection plan with equivalent number of CPs resulted in a statistically significant improvement over manual plans (paired t-test, $p < 0.001$). The described approach has potential to improve patient outcomes by reducing loss of healthy bone in tumor surgery while offering a surgeon multiple resection plan options.

Keywords — bone tumors, surgical planning, automated planning, orthopedic oncology

1 INTRODUCTION

Decisions regarding the approach of a bone tumor resection surgery are based on factors such as tumor location, size, proximity to vital structures, and the possibility of limb salvage or reconstruction (Henshaw and Malawer, 2006). The primary aim of tumor surgery is to excise the whole mass with negative margins, ensuring no tumor cells remain in the operative field. A standard practice is to excise the tumor *en-bloc* (in one piece) with a surrounding cuff of normal tissue, which reduces the likelihood of local recurrence.

Computer-assisted surgery has increased the accuracy of bone tumor resections and aids in preserving anatomical structure and function (Jaffe et al., 2013; Wong and Kumta, 2014, 2013), however planning and performing resections is a technically demanding and time-consuming process (Avedian et al., 2010; Sternheim et al., 2021; Young et al., 2015). Cutting planes (CPs) must be carefully positioned to avoid intra-lesional cuts and still allow the tumor to be resected *en-bloc* with negative margins, all while avoiding critical anatomy. Despite the significant advantages of modern planning and three-dimensional (3D) visualization software (Wong and Kumta, 2014), the placement of CPs is a manual task, often dependent on a surgeon's experience with a previous similar case.

A good solution in tumor surgery is one which both completely resects the bone tumor with no diseased tissue remaining in the host bone, and is likely to provide a good post-operative outcome and quality of life to the patient. Research has been performed on the accuracy of 3D virtual planning for multi-planar oncologic osteotomies, but detail of the methods employed in positioning of CPs in these works is limited. Aponte-Tinao et al. (2013) performed multi-planar osteotomies of tumors within the knee, where "osteotomies were planned to resect the tumor" using 3D software. Ritacco et al. (2013)

planned uni-, bi-, or multi-planar osteotomies in a 3D virtual scenario according to tumor extension, with the type of osteotomy planned determined by the tumor nature, location, and definition of clear margins. Young et al. (2015) planned and performed 18 tumor resections, with nine located around the pelvis, and seven classed as diaphyseal resections. The researchers noted that the planning phase of each case took a mean of 45 minutes (20 to 95), with the time decreasing markedly after the fifth case to a mean of 25 minutes (10 to 45). While the outcome of all these works were favorable to their area of research and were successful in treating the patient's pathologies, the only measure of the quality of the placement of CPs during pre-surgical planning was that no patients experienced a local recurrence or metastasis at final follow-up.

Automated generation of preoperative plans for orthopedic surgeries is a developing field. Carrillo et al. (2020) proposed an optimization framework to generate preoperative plans for forearm osteotomies, using a genetic algorithm capable of solving multi-objective optimization problems. The complete planning solution includes the configuration of the osteotomy plane, fixation plate, and screws, resulting in feasible clinical solutions with improvements to alignment over the gold standard.

Zhang et al. (2019) proposed a semi-automated approach to generating a bone tumor resection plan (RP), where a 'danger region' is established based on minimum safe margins as defined by a surgeon. A two-dimensional (2D) projection of the danger region is generated, with CPs set against the curve of the danger region. An algorithm first minimizes the area of empty space between the CPs and the danger region, then separately optimizes the alignment of CPs to minimize the volume of healthy bone collaterally resected with the tumor. This work may be the first to attempt to optimize the placement of CPs for bone tumor resections, however initially optimizing the RPs in 2D may generate 3D solutions optimized to a local minima, rather than a global 3D minima. Moreover, while only a single candidate RP was generated for each of the nine tumor cases in the study, the mean time for a 3D solution to be optimized was 116 minutes.

The goal of this research was the development of an algorithm capable of optimizing the position of CPs for *en-bloc* resection of bone tumors, such that healthy bone collaterally resected with the tumor is minimized, while still adhering to principles of wide-margin tumor surgery. The proposed hypotheses were that a greater number of CPs would decrease the amount of healthy bone resected with the tumor. Additionally, it was hypothesized that although a greater number of planes would result in less healthy bone collaterally resected with the tumor, there would be a diminishing return on this value. Finally, it was hypothesized that an algorithm-generated RP would save more healthy bone than a manually prepared RP with an equivalent number of cuts.

2 MATERIALS & METHOD

The proposed approach automates the generation of RPs by rotating 3D bone and tumor data about the *z* and *y* axes of Cartesian coordinate space, then positioning a CP against the tumor surface perpendicular to the *x* axis. For a given RP, each CP corresponds to a separate rotation of the bone and tumor data, with the rotation of all CPs optimized for the minimum volume of healthy bone collaterally resected with the tumor.

Anonymized image data for 20 tumor cases was acquired as part of a retrospective study with permission from the appropriate review board (approval ID: LNR/18/SVHM/21). All cases were non-metastatic malignant primary tumors (9 chondrosarcoma & 11 osteosarcoma) located near the distal femur (12) or proximal tibia (8), visualized with either CT or MR data. All images included the condylar surfaces. The images were manually segmented by the authors using an open source image processing software (MITK, DKFZ, Germany). The boundary of the bone and tumor were individually traced in each slice of the 3D image, after which an in-built marching cubes algorithm and smoothing function converted the sets of image slices into separate models. While not exhaustive, these cases were considered as good representations of the range of possible tumors around the knee (visualizations in appendices). Bone and tumor surface mesh models were then generated and imported into Matlab (Matlab R2018a, The MathWorks Inc., USA).

A 3D grid of uniformly spaced points was generated in Matlab, spanning the height, width, and depth of the bone and tumor surface meshes. Using a ray intersection method similar to the approach described by Patil and Ravi (2005), the position of each grid point was assessed to determine whether it was inside or outside the bone and tumor surface meshes. Grid points outside the surface meshes

were discarded, resulting in a homogenous 3D voxel model for the bone and tumor (2×2×2mm cubic grid spacing, coordinates set at cube centers). Bone voxels inside the tumor surface mesh were discarded, leaving two 3D arrays of surface points for the bone and tumor, and one 3D array of healthy bone voxels, from which the collaterally resected healthy bone voxels could be compared between RPs. Finally, both surface meshes and the voxel arrays were uniformly repositioned in Cartesian coordinate space such that the centroid of the tumor surface mesh was set to the coordinate space origin (x_0, y_0, z_0), aligned with the existing coordinate system of the segmented bone and tumor scan data. This process is shown in **Fig. 1**.

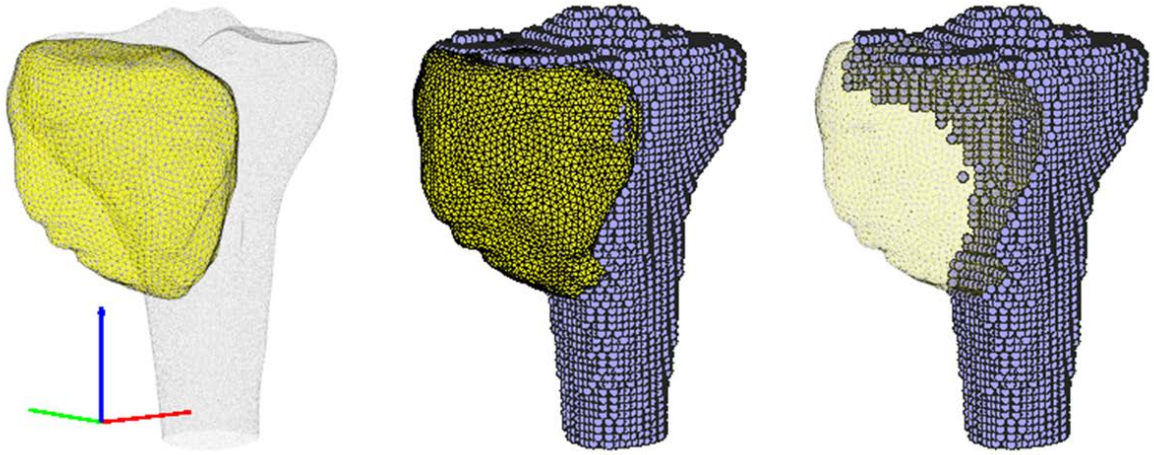

Fig. 1: Surface (left) and voxel (middle) meshes of bone and tumor with *xyz* coordinates, showing healthy bone outside tumor (right)

2.1 Positioning of Cutting Planes

For a given CP, the three arrays were uniformly rotated by two angles about the *z* and *y* axes, setting all three bone and tumor arrays in a new orientation. In this new orientation the most distant point of the tumor surface mesh in the *x* axis is set as the position of CP. All bone voxels with an *x* coordinate value less than the *x* position of the CP are marked as 'true'. After recording the true healthy bone voxels, the three arrays are reset to their initial orientation. Each additional CP adds a new pair of rotation angles, rotating the bone and tumor about the *y* and *z* axes, and setting the position of the CP as above. The total bone removed is the combination of voxels marked true by all CPs. These points are calculated by equation (1),

$$\bigcap_{i=1}^{n} n_i \cdot b_{xyz} < \max\left(n_i \cdot t_{xyz}\right) \quad (1)$$

which describes the intersection of points of resected healthy bone, where n is the number of desired CPs, n_i is the normal of CP i, b_{xyz} is the healthy bone voxels, and t_{xyz} is the tumor surface points. The values of the rotation angles which minimized bone waste were determined using particle swarm optimization.

2.2 Optimizing Cutting Plane Position via Particle Swarm

Particle Swarm Optimization (PSO) is a population-based stochastic optimization technique, capable of comparing many solutions for a single problem across a broad search space (Kennedy and Eberhart, 1995). In this work, PSO is utilized to determine the position and orientation, or pose, of a desired number of CPs. The objective function is 'bone waste', the number of true healthy bone voxels collaterally resected by all CPs, which the PSO algorithm attempts to minimize by optimizing the pose of the CPs. Using PSO, the *z* and *y* rotations for each CP are encoded in a population of multivariable particles, such that particles optimizing a RP with nCPs are encoded with $2 \times n$ CPs, with the value of each *z* and *y* rotation angle bounded between ±180°.

Calculating the objective function at every discrete *z* and *y* rotation angle (within the particle bounds) produces a value of bone waste for every CP pose, visualizing the search space and the regions of minima. For a single CP, the encoding particles will converge toward regions of minima, and the particle with the best objective function sets the pose of the CP with its encoding values. An example of this process is shown in **Fig. 2**. The blue and yellow regions are the value of the objective function across the entire search space, i.e., the bone waste at every *z* and *y* rotation angle combination. In the figure, the red mark indicates the objective function for a particular CP pose.

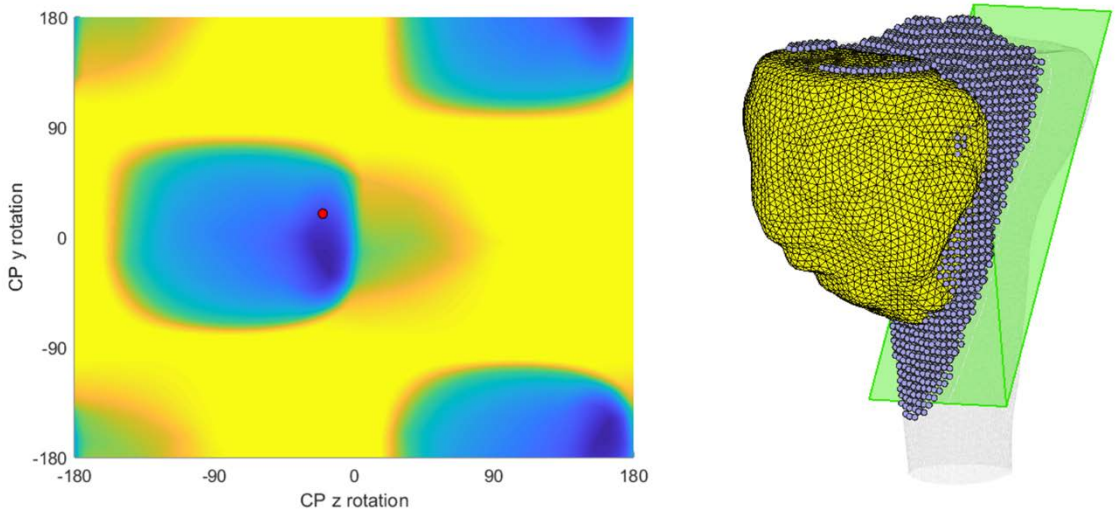

Fig. 2: An example of an optimized encoding particle in the *z* and *y* search space (left), where the dark blue regions indicate minimum bone waste, and yellow indicates high bone waste, and the corresponding optimized CP set against the tumor and collaterally resected healthy bone (right).

2.3 Refinement of search space

The position of a particle in the search space is equivalent to the combined pairs of *z* and *y* rotation values for each CP. Every additional CP increases the dimensionality of the search space by two, with the number of candidate solutions in the search space for n CPs increasing by a power of $2 \times n$. Without a significant population of particles in the search space, a swarm is less likely to converge on the global minima as n increases.

A method of progressively narrowing the search space would be advantageous for efficient optimization of multiple CPs. Thus, a single-to-multi-planar refinement approach was investigated, where the n^{th} CP in a RP is optimized individually, then all CPs are optimized simultaneously. Bounds are applied to the *z* and *y* rotations of existing CPs, based on the corresponding values of the best $n\text{-}1^{th}$ RP, plus the best values for the newly added n^{th} CP. These bounds are incrementally reduced after successive optimizations, decreasing the breadth of the search space to target the region containing the global minima.

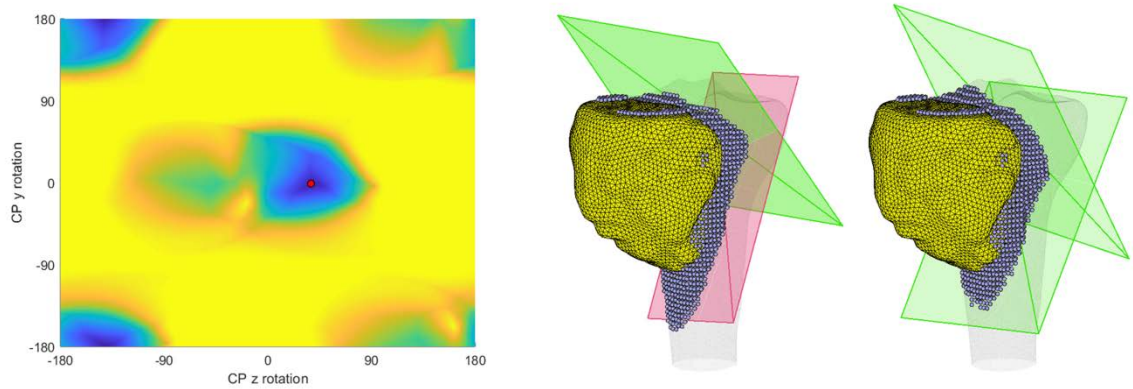

Fig. 3: The search space of the 2^{nd} CP (left), with the 1^{st} CP fixed in position (middle), and the final optimized pose (right). Note the difference in position of the fixed CP (red) compared to its free (green) configuration.

For each tumor case, the optimization process began with a single CP (RP1). After 20 individual optimizations of RP1, the pose of the 1st CP was fixed to the best result. This configuration and the corresponding waste bone is the input to the single-planar refinement stage of optimizing two CPs (RP2), with one fixed CP, and one unconstrained CP (**Fig. 3**). With the 1st CP fixed to the best configuration of RP1, the 2nd CP moves freely about the search space. After optimizing the 2nd CP five times, the best RP is the input to the multi-planar refinement stage. The bounds on the particle values are restricted to ±60° of the best RP, and all CPs are optimized within the bounded search space. This narrowing of the particle value bounds allows 120° of movement for each CP, allowing the swarm to shift the search space bounds after successive optimizations, and converge toward a distant region of minima. After another five optimizations, the search space bounds is narrowed further to ±30°, then finally the bounds of the last five optimizations are restricted to ±15° of the best RP values generated thus far.

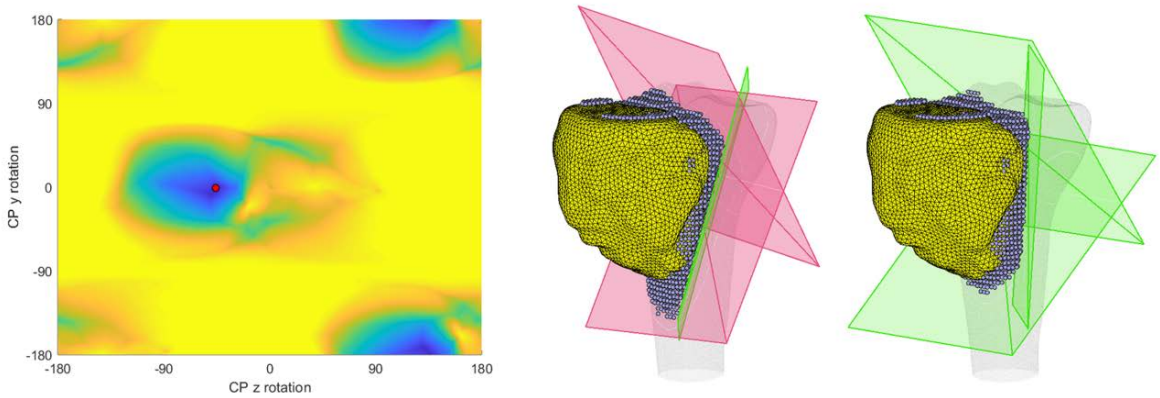

Fig. 4: The search space of the third CP (left), with the two existing CPs (red) fixed in position (middle), and the final optimized resection (right). Note the pose of the fixed CPs compared to the optimized RP in Fig. 3

Generalizing, the optimization of RPn begins with the CPs fixed in the pose of RPn-1 **(Fig. 4)**. Two new values encode the nth CP, which is free to move about the ±180° search space. After five optimizations, all CPs can move within ±60° of the current best, narrowing to ±30° for the next five optimizations, and ±15° for the final five optimizations. This process is then repeated for RPn+1 **(Fig. 5)**. The number of optimizations, the incremental restriction to the search space, and the number of particles in the swarm were chosen based on initial experimentation, and the knowledge that the stochastic nature of PSO cannot guarantee convergence upon a global minima. The selection of the values restricting the bounds of the multi-planar stages of the optimization was a trade-off between substantially increasing the effective swarm density; providing the CPs a broad search space and wide range of motion; eliminating regions of search space that were unlikely to contain the global minima based on the results of the single-planar stage; and reducing the overall optimization time. It was considered unlikely that, after optimizing a CP in the single-planar stage, converging on the global minima would require previously optimized CPs to move to the opposite side of the tumor.

The PSO parameters utilized were default Matlab values, with both the personal and global learning coefficients set to 1.49; an adaptive inertia based on the values of the personal and global learning coefficients for each particle, with a random initial weight for each particle between 0.1 and 1.1; and a maximum number of iterations of 200 × (nCPs × 2), with an early termination condition of 20 iterations without change.

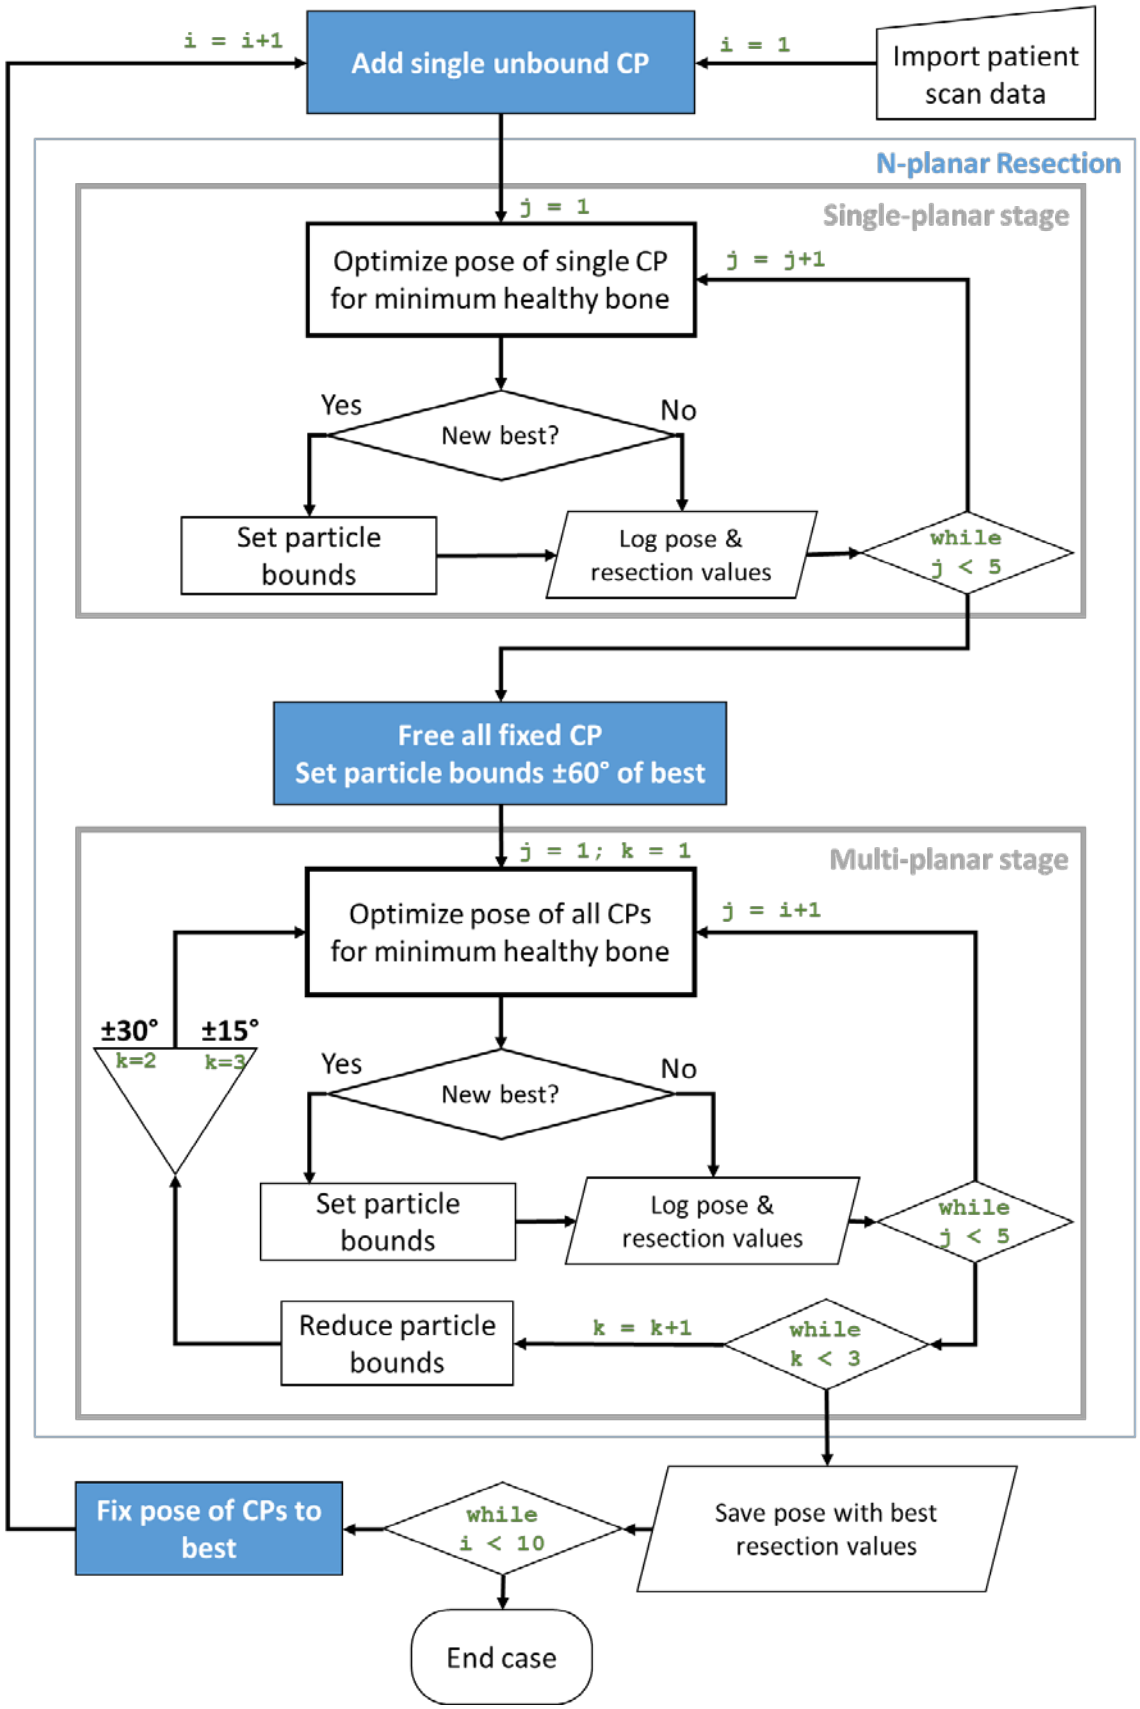

Fig. 5: System diagram of single-to-multi-planar CP optimization. Each case is optimized starting from a single CP. CPs are added to the patient data one at a time, and individually optimized in the single-planar refinement stage. The existing CPs are then simultaneously optimized along with the new CP in the multi-planar refinement stage.

2.4 Geometric Validation of Proposed Resection Plans

For four or more CPs, it is possible that the optimal RP geometry may be aligned in such a way that it cannot be removed from the bone. To ensure only valid RPs are generated by the PSO (i.e., can be removed from the bone), the RP is assessed for an 'exit path'.

This assessment of viability is performed on every particle encoding a candidate RP. A viewing angle is set as the line of intersection between a pair of CPs, then with respect to the selected viewing angle, the relative alignment of each CP is determined. For a single particle, if any CP is found to occlude the resection from that particular viewing angle, the algorithm realigns the viewing angle with the next pair of CPs, and checks whether any CP occludes the resection with respect to the new view (**Fig. 6**). If a particular view results in each CP having a negative relative *z* normal value, then no CP occludes the resection, and the particle's candidate RP is geometrically viable. Conversely, if no pair of CPs provides a view which allows removal, then that particle's candidate RP is penalized, and the algorithm checks the candidate RP of the next particle in the swarm.

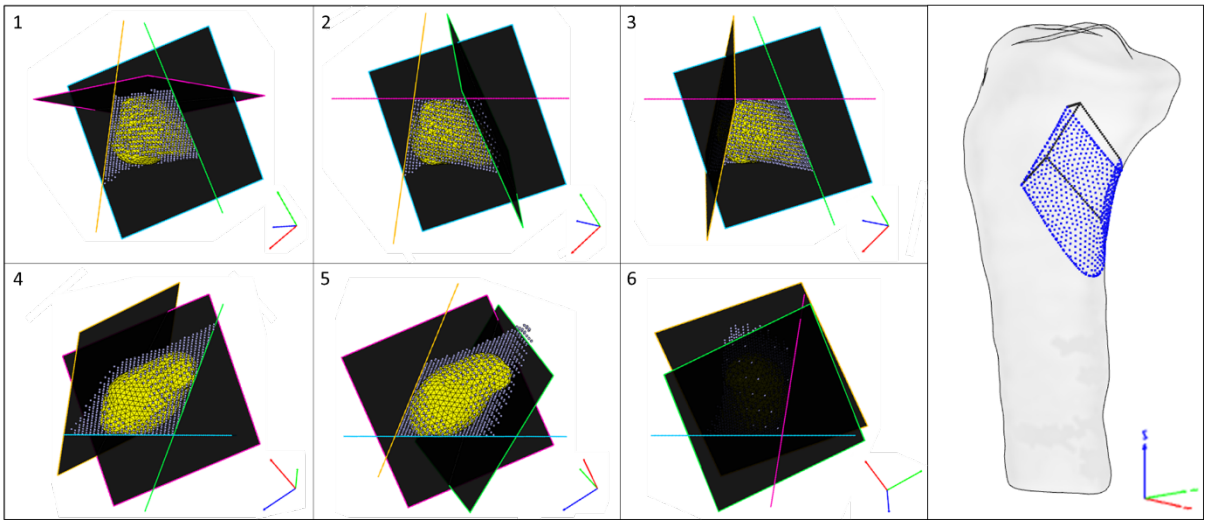

Fig. 6: Six aligned projections of a proposed 4-planar resection, and the resulting implant window. Each view is aligned with the line of intersection between a pair of CPs. CPs are colored yellow, green, cyan, and magenta. Red-green-blue axes respectively indicate the base *xyz* coordinate axes. Views 2, 3 and 6 show that the RP cannot be removed along the corresponding view, as the green and yellow CPs are slightly occluding the resections in view 2 and 3 respectively, while in view 6 the green CP is fully occluding the resection. Views 1, 4 and 5 show alignments in which the resection can be removed, as no CP occludes the resection by pointing 'toward' the viewing angle.

2.5 Manually Planned Resections

In collaboration with an experienced surgeon, a manual RP was generated for all 20 cases in a custom resection planning software. This software allowed visualization of the bone and tumor structures and placement of planes at a given distance from the tumor with the positions and orientations of planes stored for later analysis. To match the algorithm, the surgeon was instructed to plan the cases assuming that the tumor surface mesh included a cuff of healthy tissue and comprised everything that must be removed to adhere to the principles of wide-margin tumor resection surgery. Additionally the surgeon was informed all cases would have intraoperative computer navigation available in surgery, and the intended reconstruction could be to their choosing.

3 RESULTS

RPs with 1 to 10 CPs were generated for each of the 20 test cases according to the method described above. Analysis of results was performed using Matlab, and Excel (Excel 2013, Microsoft, USA). Results are reported as bone waste: the healthy bone removed as a percentage of the total resection volume, where the total resection volume is the volume of resected bone plus the volume of the tumor within the bone. **Fig. 7** compares the algorithm RPs to those manually prepared by the surgeon, in sequential case order, in groups of five. For each case the manual RP is indicated by a triangle of the

same color. For ten cases (1, 2, 3, 5, 6, 9, 11, 12, 18, 20) the algorithm generated a RP with less bone waste than the manual RP using one fewer CP. For four cases (10, 13, 14, 19) the algorithm generated a RP with less bone waste than the manual RP using two fewer CPs, and for case 15, a RP was generated with less bone waste than the manual using three fewer CPs. The volume of bone waste for the remaining five cases (4, 7, 8, 16, 17) was a similar value for both the manual and algorithm RP.

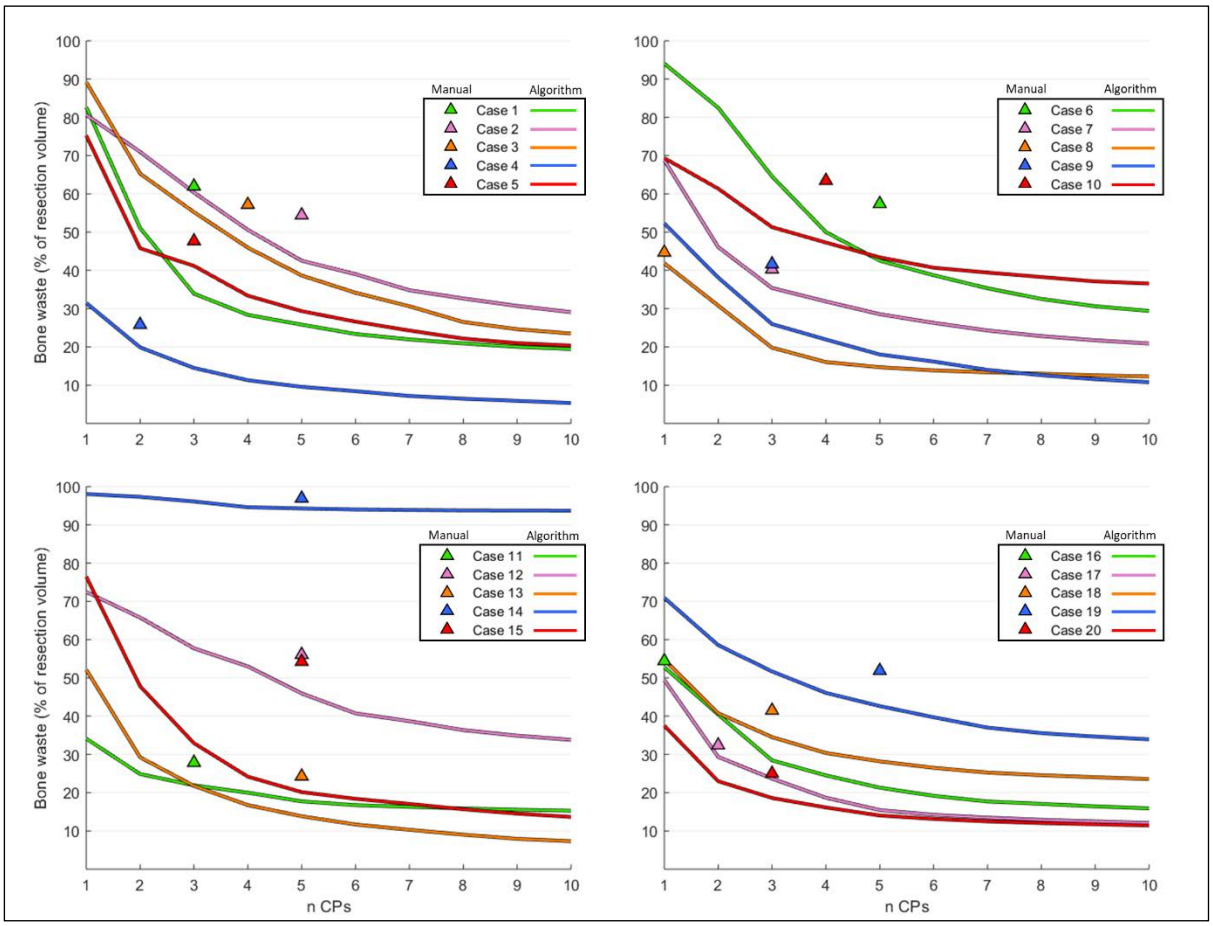

Fig. 7: Comparison of bone waste between manual and algorithm-generated RPs

Fig. 8 compares the percentage change in absolute waste between the best n-planar RP and the best n_{+1}-planar RP, i.e., the rate at which absolute bone waste changes with each additional CP. Considering the improvement to absolute bone waste with each additional CP; All 20 cases improved by over 20% with the 2nd CP; 18 of 20 improved by over 20% with the 3rd CP (over 10% for remaining two), 12 of 20 improved by over 20% with the 4th CP (over 10% for other eight cases), and with the 5th CP, only eight of the 20 cases improved by over 20%, while ten improved by between 10 and 20%, and the remaining two cases improved by less than 10%. Overall, there are diminishing returns on minimization of bone waste using 6 or more CPs, with the majority of the improvements provided using up to 5 CPs.

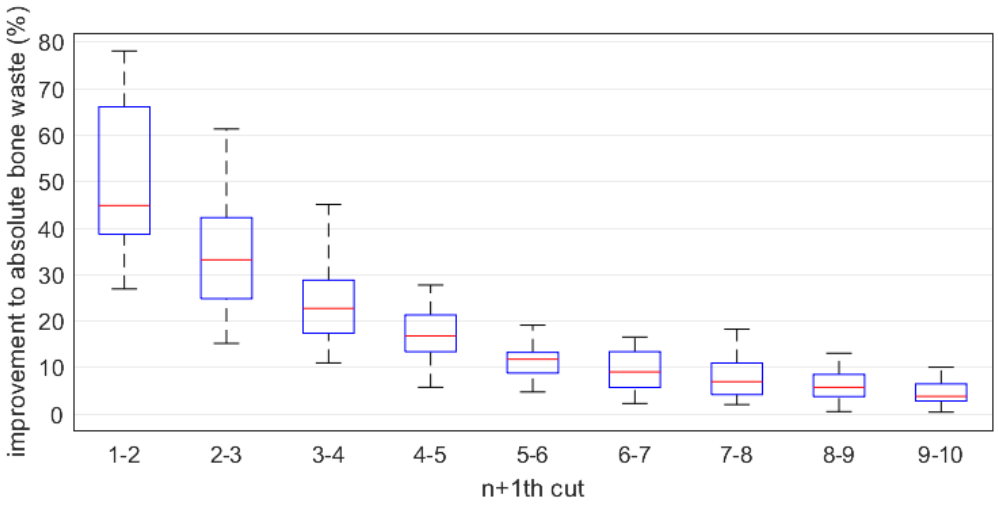

Fig. 8: Average change to absolute bone waste with each additional CP

Fig. 9 shows the percentage improvement to absolute waste between the manual RP and the algorithm RP using the equivalent number of CPs. For all tumor cases, on the metric of minimizing bone waste, the RPs prepared by the surgeon were outperformed by the algorithm for the equivalent number of CPs. The greatest difference was on case 15, with an improvement to absolute bone waste of 78.72%, and the minimum on case 16, with an improvement of 6.42%. The algorithm equivalent RPs resulted in a statistically significant improvement over manual (paired t-test, $p < 0.001$), with a mean improvement of 32.44%.

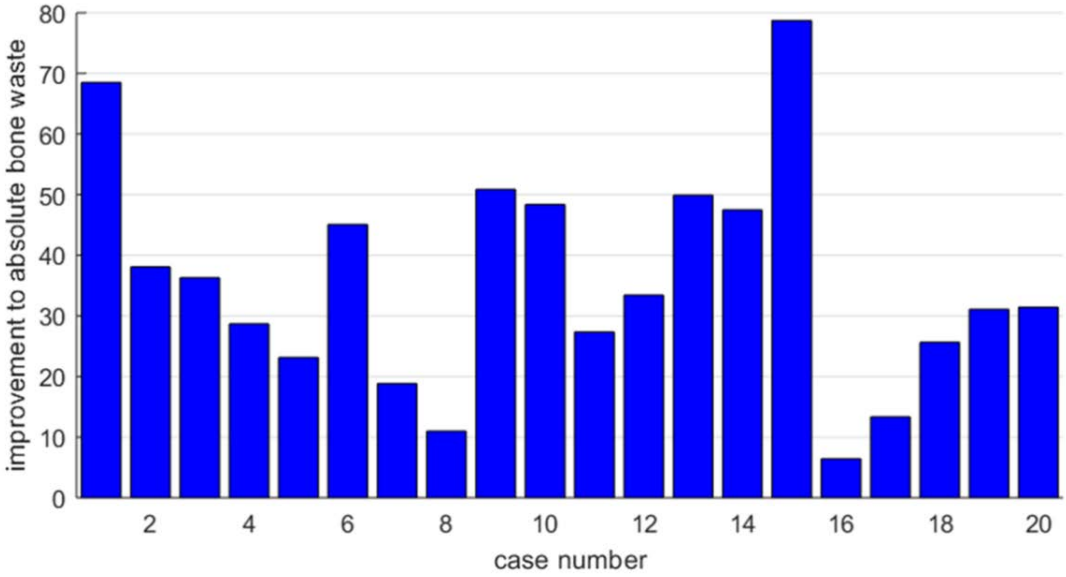

Fig. 9: Improvement to absolute bone waste from manual to algorithm-equivalent RPs

Fig. 10 shows the time taken for the algorithm to optimize n CPs for each case. The value is cumulative, meaning the time taken to optimize RPs with nCPs includes every previous RP. RPs with fewer CPs are generated much faster than those using 6 or more CPs. The slowest case to optimize all 10 CPs was case 18, in 260 minutes, while case 20 was the fastest, performed in 114 minutes. Considering RPs using six or fewer CPs, case 18 was again slowest to optimize at 84 minutes, while optimization of case 20 was completed in 39 minutes. **Fig. 11** shows the total optimization time relative to the total number of surface and voxel points in each case, indicating cases with a greater number of points are slower to optimize.

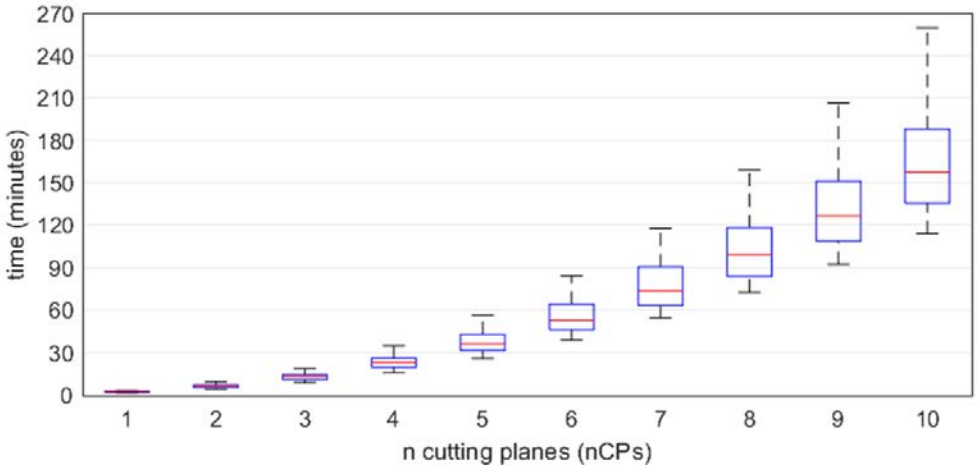

Fig. 10: Total time taken for algorithm to generate RPs with n CPs

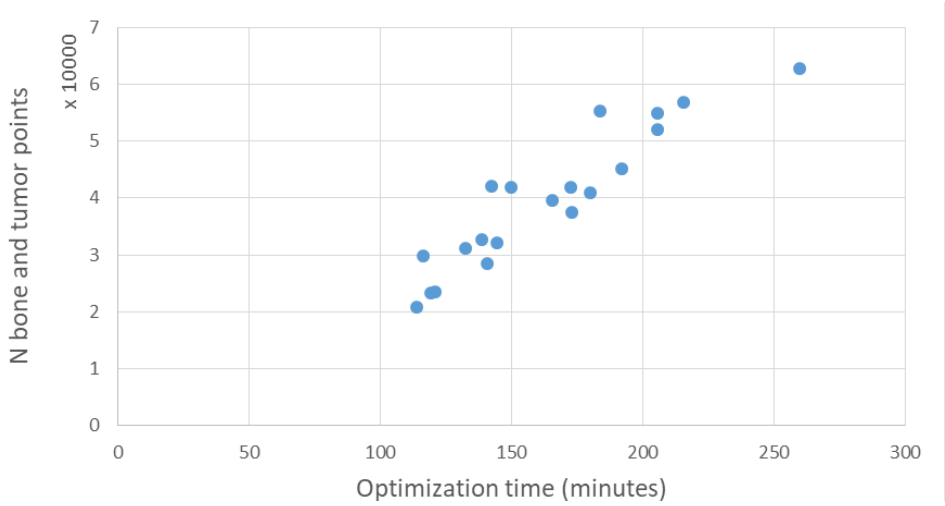

Fig. 11: Total optimization time and number of surface and voxel points

Fig. 12 shows case 14, a tumor with a higher bone waste due to the lower cortical infiltration of the bone, and the concave nature of the tumor. Additional examples of RPs are shown in **Fig. 13**.

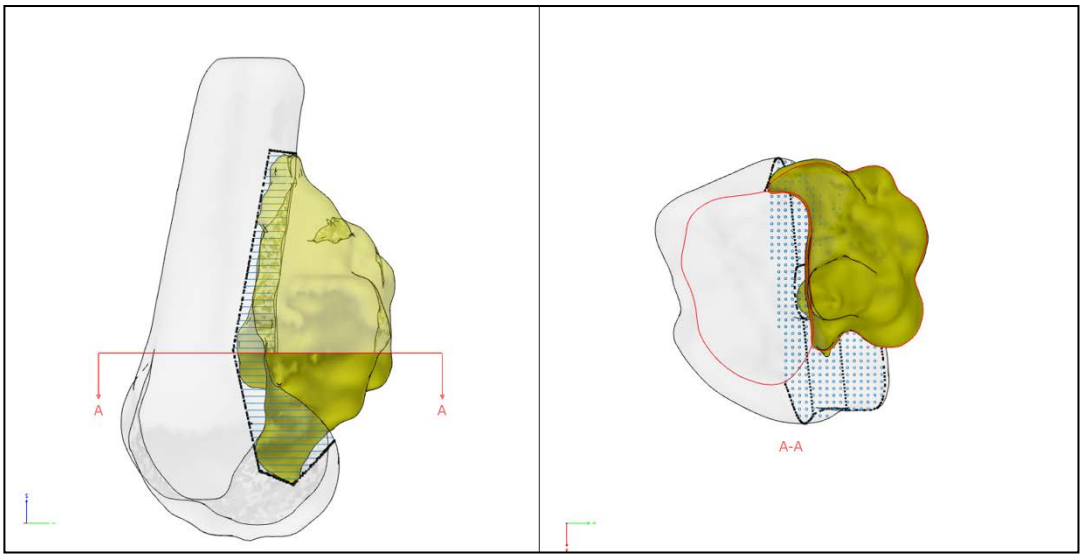

Fig. 12: Cross-sectional view of case 14 and its manual resection plan

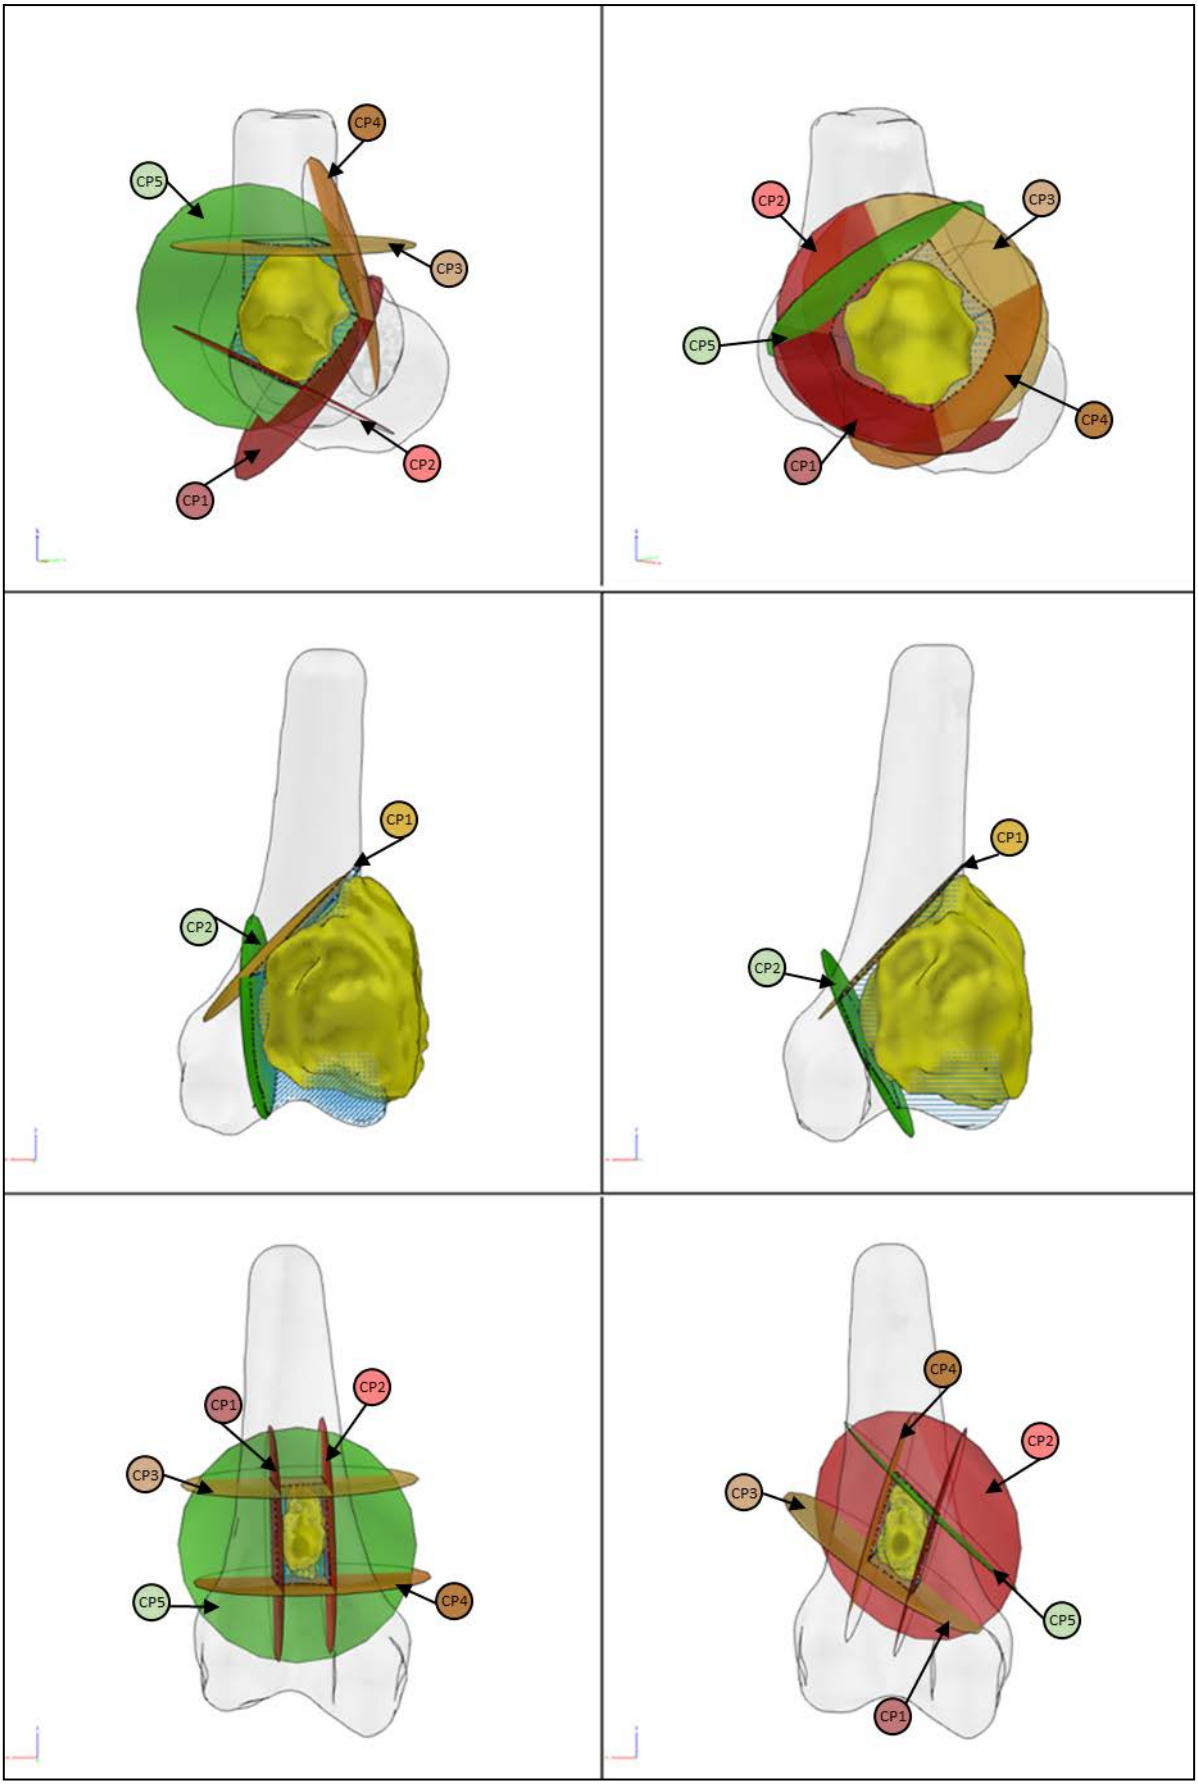

Fig. 13: Examples of manual (left) and algorithm (right) RPs for case 15, case 17, and case 6, with CPs marked by number and color.

4 DISCUSSION

This work has described an approach for automated optimized placement of CPs in orthopedic oncology surgery, minimizing the amount of healthy bone removed with the tumor. The described approach efficiently configures the placement of multiple CPs, while ensuring the resulting geometry can be removed from the bone. From 20 patient cases, it was demonstrated that a greater number of CPs results in a reduction to the volume of healthy bone collaterally resected with the tumor, and that this reduction exhibits diminishing returns with 6 or more CPs. The method performed the optimization in a timely manner and saved a greater volume of healthy bone than equivalent manually prepared RPs, with no manual user input beyond segmentation of the bone and tumor. Representing the volume of the bone with voxels enable quick Boolean calculation of which points of bone are inside or outside the RP, and thus allow the optimization algorithm to quickly minimize the resected voxels, and therefore the RP volume. Inclusion of the bone surface points is not necessary.

This research innovates upon the previously described semi-automated approach from Zhang et al., as the method can be used to optimize a resection with any number of cutting planes, requiring no manual interaction other than input of data. The method also ensures all resection plans improve as more CPs are added, and the results of the approach demonstrate algorithm-generated RPs are at least as good as the manual approach, with respect to bone waste.

The Gold Standard measurement for quality of a resection is against a manual plan. Manually positioning CPs in a virtual environment is the current method of generating a resection plan, as there is not yet a reliable alternative approach. The inherent unique nature of bone tumor cases increase the difficulty of comparing the results of different planning methodologies, meaning measurements can currently only be made against a manually planned resection for the same case. Due to the relative infancy of the field, there is not yet a consistent value or metric assigned to this Gold Standard measurement.

In this work, the results quantify a resection plan by its efficiency; that being the percentage of bone waste compared to the intra-cortical volume of the tumor.

Previous research has assessed resection plans generated by an automated method to one manually prepared by a surgeon, comparing the total volume of the resected specimen from each. A raw measurement of volume provides no indication of how much of that resected specimen must be removed, how much more could have been preserved, or the composition of tissues.

While there may be clinical significance to the exact volume of a tumor, from a standpoint of resection efficiency, the volume of the tumor itself is somewhat irrelevant; no matter its size, it must be removed. Furthermore, in instances of tumors with large extra-cortical extension, comparing the quantity of bone waste to the entire tumor volume gives a false improvement to the efficiency of the plan.

Reporting the bone waste as a percentage of intra-cortical tumor volume, rather than a raw measurement of the volume of the resection plan, negates the volumetric variations between small and large tumors, which in turn allows the identification and analysis of trends and commonalities across all tumor cases. Reporting results in this manner may also enable analysis or measurements between patients and assist in understanding clinical aspects, such as tumor infiltration compared to volume.

4.1 Effectiveness of algorithm

As an assistive tool for tumor surgery planning, the algorithm generates RPs that save more healthy bone than what can be achieved manually as demonstrated on 20 example tumor cases.

The single-to-multi-planar search space refinement ensured every n^{th} RP improved upon the n^{th}-1 RP. For RPs with 8 or more CPs, in only five cases was the multi-planar refinement stage unable to improve upon the best RP generated by the single-planar refinement stage. All other cases had at least one best RP with 8 or more CPs from the multi-planar refinement stage.

With regard to the apparent outlier in **Fig. 7**, case 14 is a large concave surface tumor with minimal infiltration. As the measure of bone waste is taken with respect to the number of tumor voxels within the bone surface mesh, tumors with a large extra-cortical portion are effectively reduced in size, as the tumor surface mesh is still used to calculate CP position. While this gives the appearance of the algorithm underperforming on tumors with minimal infiltration, the improvement to bone waste with each additional CP (**Fig. 8**) is comparable to other tumor cases. Thus, assessing the effectiveness of

the algorithm for a given case cannot be made solely from the bone waste at every CP, and should instead consider the improvement to bone waste from each CP.

As compared to existing studies, the work by Zhang et al. generated nine individual RPs semi-autonomously in a mean time of 116 minutes, with a surgeon specifying the number of CPs required. Young et al. reported the RPs for the first five cases took a mean of 45 minutes (20 to 95) to plan, and decreased thereafter for the remaining 13 cases. In this research, the total optimization time for each case can be attributed to the total number of points being manipulated in the algorithm (**Fig. 11**). For case 20, the total number of points from the tumor and bone surface meshes and the healthy bone voxels was 20,781, while case 18 contains 62,757 points. All cases were optimized for at least 6 CPs in a mean time of 55 minutes (39 to 83). This shows the algorithm optimizes RPs significantly faster than the current best for automated generation of RPs. While this also shows the algorithm as taking an additional 10 minutes over the manual approach, the algorithm produces six RPs from which a surgeon can select, whereas the stated automated method and standard manual approach produce only one resection per case.

Of the nine RPs generated by Zhang et al., two had a greater volume than the manual approach. In this study, all algorithm-generated RPs improved upon the manual with an equivalent number of CPs, and in some cases provided a volumetrically equivalent RP with fewer CPs.

The time taken to generate the manual RPs for this study was not included, as the surgeon was familiar with many of the cases. Furthermore, the custom software used to generate the manual RPs was designed for ease of CP placement, and is not a true representation of the time required or difficulty associated with planning a case with clinical software.

4.2 Qualitative assessment of algorithm RPs

Reviewing 3D models (**Fig. 13, appendices**) of all best RPs, the geometric validation check successfully ensured every resection could be removed from the bone, however in a number of cases, blind cuts would be required to separate the resection mass from the remaining healthy bone. For all cases, no blind cuts were required when using 4 or fewer CPs, and for 17 cases (excluding cases 1, 4, 6), no blind cuts were required when using 5 CPs. Across all 200 optimized RPs, 24 used more than one blind cut, and in only one RP (case 6, 6 CPs) did multiple blind cuts occur in a RP with fewer than 7 CPS. It should be noted that a single blind cut was used in two of the manual RPs, meaning a blind cut is not necessarily disqualifying for the viability of a RP, but is not a common approach.

A shortcoming of the algorithm exists in the generation of single-planar RPs, in which the bone surface mesh is only a partial scan of the patient's limb. As the algorithm will attempt to minimize bone waste, there are instances where the best single-planar RP preserves bone around the joint surface, and instead resects all healthy bone distal from the affected joint. An example of this is in case 8, where the manual RP involved a single plane placed orthogonal to the diaphysis of the femur, proximal to the tumor, resulting in either an amputation, or a mega-prosthesis reconstruction. The algorithm equivalent preserves a single condyle, resecting all remaining bone distal from the joint, effectively removing the entire femur. This would be prevented with a full scan of the patient's bone, as there would be sufficient healthy bone distal of the affected joint for the algorithm to more correctly place the CP.

3D representations of all cases and RPs are available to be viewed in the **supplementary material**.

4.3 Future work

The approach is currently implemented in Matlab, and is not optimized for speed. Improvements could be made to processing time by removing many of the visualization steps, or performing the calculations analytically in another programming language. Furthermore, applying the algorithm to new patient cases would allow for an accurate measure of the difference in time taken between standard manual planning and automated generation of a RP.

The approach described in this work has only been tested on bone tumors around the knee, however the algorithm is agnostic to the affected bone and the location of the tumor. Opportunities for future research exist in the application to tumors in other locations around the body, and the benefits that may be provided.

In this study, all cases were monolithic single tumor masses. If the tumor were fragmented, the algorithm could either treat the mass as a single non-contiguous tumor, or as multiple individual tumors,

with each requiring its own optimized RP. Treating the fragments as a single tumor would result in a high amount of collaterally resected bone relative to the tumor volume, but the method would perform as described. Treating each fragment separately should reduce bone waste compared to monolithic approach, but would take longer to optimize. Additionally, if the fragments are in close proximity, the 'remaining healthy bone' for the first fragment would be an input to the second fragment, and a globally optimized solution may depend on the order in which the fragments are optimized.

While the margin distance in this study was set to zero, the algorithm allows for easy alteration of this value by offsetting the plane position from the tumor surface. Alternatively, the size of the tumor can be scaled in 3D, generating an over-sized tumor equivalent to the safety margin.

The algorithm optimizes RPs by minimizing the number of collaterally resected healthy bone voxels, with all bone voxels weighted equally when measuring the resection volume. Using more precise segmented models would allow a different weighting to be applied to cortical or trabecular bone. Additionally, any number of meshes can be passed as inputs to the PSO process, with all points comprising those meshes undergoing the same Boolean intersection check described in equation (1). Assigning a weighting or penalty value to these points would make it possible to incentivize RPs which avoid specific critical anatomy, such as joint surface or ligament attachment points. Future research will examine the algorithms' capability of optimizing RPs with additional segmented critical anatomy included as tertiary inputs.

While no tool type has been specified, the research assumes the RPs would be performed with an infinitely thin tool capable of planar cuts, likely a saw or burr, with the assistance of either computer navigation or a surgical robot. As the findings are agnostic to tool selection, future research would include identifying which tool has optimal performance on metrics such as bone loss due to tool kerf, accessibility for all CPs, or whether certain resections necessitates the development of novel cutting tool designs.

Most significantly, the research does not consider the viability of the resections from a clinical standpoint. Assessment of clinical viability for each RP would require consideration of factors such as biomechanics of remaining bone and reconstructive implant through finite element analysis; soft-tissue access requirements; implant fixation methods; interference from the RP with local critical anatomy such as joint surfaces; likelihood of remaining bone experiencing avascular necrosis; complexity of the RP and the benefit to the patient, compared to manual RPs, and; changes to surgical duration and the impact on intraoperative factors including infection risk or blood loss. Research investigating these clinical viability factors compared to the existing results will be the subject of future work.

5 CONCLUSION

This work presents a novel method of automating the generation of surgical resection plans for bone tumors, whereby the volume of healthy bone collaterally resected with the tumor is minimized. To the authors' knowledge, this is the first research describing a fully automated method of generating a bone tumor resection plan. The described approach efficiently configures the placement of multiple cutting planes, while ensuring the resulting geometry can be removed from the bone. From 20 patient cases, it was demonstrated that a greater number of cutting planes results in a reduction to the volume of healthy bone collaterally resected with the tumor, and that this reduction exhibits diminishing returns beyond five cuts. The method performed the optimization in a timely manner and saved a greater volume of healthy bone than equivalent manually planned resections, with no manual user input beyond segmentation of the tumor.

6 ACKNOWLEDGEMENTS

The authors acknowledge the facilities, and the scientific and technical assistance of the RMIT Advanced Manufacturing Precinct, and thank Ulrich Buehner for his contributions to the manuscript.

7 FUNDING

This work was supported by the Innovative Manufacturing Cooperative Research Centre (IMCRC/STR/180917).

8 REFERENCES

Aponte-Tinao, L.A., Ritacco, L.E., Ayerza, M.A., Muscolo, D.L., Farfalli, G.L., 2013. Multiplanar Osteotomies Guided by Navigation in Chondrosarcoma of the Knee. Orthopedics 36, e325–e330. https://doi.org/10.3928/01477447-20130222-21

Avedian, R.S., Haydon, R.C., Peabody, T.D., 2010. Multiplanar osteotomy with limited wide margins: A tissue preserving surgical technique for high-grade bone sarcomas. Clin. Orthop. Relat. Res. 468, 2754–2764. https://doi.org/10.1007/s11999-010-1362-0

Carrillo, F., Roner, S., von Atzigen, M., Schweizer, A., Nagy, L., Vlachopoulos, L., Snedeker, J.G., Fürnstahl, P., 2020. An automatic genetic algorithm framework for the optimization of three-dimensional surgical plans of forearm corrective osteotomies. Med. Image Anal. 60, 101598. https://doi.org/10.1016/j.media.2019.101598

Henshaw, R., Malawer, M., 2006. Review of Endoprosthetic Reconstruction in Limb-sparing Surgery. Musculoskelet. Cancer Surg. 383–403. https://doi.org/10.1007/0-306-48407-2_25

Jaffe, N., Puri, A., Gelderblom, H., 2013. Osteosarcoma: evolution of treatment paradigms. Sarcoma 2013, 203531. https://doi.org/10.1155/2013/203531

Kennedy, J., Eberhart, R., 1995. Particle swarm optimization, in: Proceedings of ICNN'95 - International Conference on Neural Networks. IEEE, pp. 1942–1948. https://doi.org/10.1109/ICNN.1995.488968

Patil, S., Ravi, B., 2005. Voxel-based representation, display and thickness analysis of intricate shapes. Proc. - Ninth Int. Conf. Comput. Aided Des. Comput. Graph. CAD/CG 2005 2005, 415–420. https://doi.org/10.1109/CAD-CG.2005.86

Ritacco, L.E., Milano, F.E., Farfalli, G.L., Ayerza, M.A., Muscolo, D.L., Aponte-Tinao, L.A., 2013. Accuracy of 3-D Planning and Navigation in Bone Tumor Resection. Orthopedics 36, e942–e950. https://doi.org/10.3928/01477447-20130624-27

Sternheim, A., Rotman, D., Nayak, P., Arkhangorodsky, M., Daly, M.J., Irish, J.C., Ferguson, P.C., Wunder, J.S., 2021. Computer-assisted surgical planning of complex bone tumor resections improves negative margin outcomes in a sawbones model. Int. J. Comput. Assist. Radiol. Surg. 16, 695–701. https://doi.org/10.1007/s11548-021-02337-w

Wong, K.C., Kumta, S.M., 2014. Use of Computer Navigation in Orthopedic Oncology. Curr. Surg. Reports 2, 47. https://doi.org/10.1007/s40137-014-0047-0

Wong, K.C., Kumta, S.M., 2013. Computer-assisted Tumor Surgery in Malignant Bone Tumors. Clin. Orthop. Relat. Res. 471, 750–761. https://doi.org/10.1007/s11999-012-2557-3

Young, P.S., Bell, S.W., Mahendra, A., 2015. The evolving role of computer-assisted navigation in musculoskeletal oncology. Bone Jt. J. 97-B, 258–264. https://doi.org/10.1302/0301-620X.97B2.34461

Zhang, Y., Li, F., Qiu, L., Xu, L., Niu, X., Sui, Y., Zhang, S., Zhang, Q., Zhang, L., 2019. Toward Precise Osteotomies: A Coarse-to-fine 3D Cut Plane Planning Method for Image-Guided Pelvis Tumor Resection Surgery. IEEE Trans. Med. Imaging 14, 1–1. https://doi.org/10.1109/tmi.2019.2951838

9 APPENDICES

Supplementary material available through online publication