See, Plan, Cut: MPC-Based Autonomous Volumetric Robotic Laser Surgery with OCT Guidance

Ravi Prakash [1*†], Vincent Y. Wang [1*], Arpit Mishra [1], Devi Yuliarti [1], Pei Zhong [1],
Ryan P. McNabb [2], Patrick J. Codd[1,3‡], Leila J. Bridgeman[1‡]

Abstract—**Robotic laser systems offer the potential for sub-millimeter, non-contact, high-precision tissue resection, yet existing platforms lack volumetric planning and intraoperative feedback. We present RATS (Robot-Assisted Tissue Surgery), an intelligent opto-mechanical, optical coherence tomography (OCT)-guided robotic platform designed for autonomous volumetric soft tissue resection in surgical applications. RATS integrates macro-scale RGB-D imaging, micro-scale OCT, and a fiber-coupled surgical laser, calibrated through a novel multi-stage alignment pipeline that achieves OCT-to-laser calibration accuracy of 0.161 ± 0.031 mm on tissue phantoms and *ex vivo* porcine tissue. A super-Gaussian laser–tissue interaction (LTI) model characterizes ablation crater morphology with an average RMSE of 0.231 ± 0.121 mm, outperforming Gaussian baselines. A sampling-based model predictive control (MPC) framework operates directly on OCT voxel data to generate constraint-aware resection trajectories with closed-loop feedback, achieving 0.842 mm RMSE and improving intersection-over-union agreement by 64.8% compared to feedforward execution. With OCT, RATS detects subsurface structures and modifies the planner's objective to preserve them, demonstrating clinical feasibility.**

To our knowledge, this is the first demonstration of closed-loop volumetric robotic laser resection with OCT guidance, enabling precise, obstacle-aware tissue removal with potential applications in neurosurgery and complex soft tissue oncology.

Index Terms—**Robotic Surgery, Volumetric Resection, Laser–Tissue Interaction (LTI), Optical Coherence Tomography (OCT), Model Predictive Control (MPC)**

I. INTRODUCTION

Robotic-assisted surgery has significantly improved precision, dexterity, and safety in minimally invasive procedures [1], [2], largely through the integration of novel instruments with complementary sensing and control. However, most systems still rely heavily on human teleoperation for critical tasks like tumor resection, which demand volumetric tissue removal while preserving nearby anatomical structures (e.g., vessels, nerves). These challenges are especially pronounced in neurosurgery, where tumor margins are irregular and critical neurovascular structures often traverse the lesion.
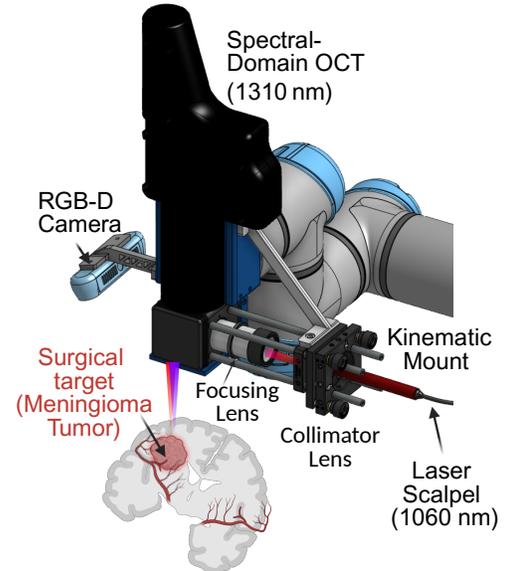
Fig. 1. **RATS Platform:** Hardware configuration for OCT-guided volumetric laser surgery: The robot-mounted platform integrates a spectral-domain OCT scanner (1310 nm) and fiber-coupled surgical laser (1060 nm), co-axially aligned with OCT through a dichroic mirror. Collimated laser scalpel beam is focused through an achromatic doublet-based focusing lens. An RGB-D camera provides macro-scale information over OCT's micro-scale features. The system is shown targeting a simulated meningioma tumor embedded within vascular tissue, representative of neurosurgical environments requiring sub-millimeter precision and constraint-aware planning [11]

Laser-based scalpels offer sub-millimeter resolution [3], non-contact energy delivery, and simultaneous ablation and coagulation [4], [5], making them ideal for precise soft tissue removal. However, their inherent precision makes them challenging to dexterously wield manually during surgeries, thus blunting their efficacy and usage due to variability. Robotic systems have demonstrated progress in laser manipulation [6], sensing [7], and control [8], for ablation tasks, but remain constrained by three major limitations: (i) lack of integrated volumetric intraoperative sensing, (ii) poor calibration between imaging and ablation paths, and (iii) absence of constraint-aware planning informed by laser–tissue interaction (LTI) dynamics. As a result, current approaches rely on open-loop raster scanning [9] or simulation-only predictive planning [10], limiting clinical utility in complex 3D surgical environments.

Recent efforts aimed to close this gap using high-resolution imaging modalities for crater reconstruction and subsurface feature detection [12], [13]. Optical coherence tomography (OCT) [14] is particularly promising, offering micron-scale resolution several millimeters deep. When paired with a 6-

DoF robotic arm and a laser resection tool, OCT enables intraoperative volumetric planning [15], [16]. However, enabling closed-loop robotic laser resection requires coordination across multiple modules, such as an accurate LTI model, a constraint-aware volumetric planner, and precise calibration between imaging and actuation to achieve accurate autonomous or semi-autonomous surgical resection. Such closed-loop integration is especially critical in neurosurgical oncology, such as in meningioma resection, where tumors may encapsulate critical vasculature and manual resections are both time-consuming and high-risk. With these challenges in mind, we introduce RATS (Robot-Assisted Tissue Surgery), a novel OCT-guided opto-mechanical robotic platform for intelligent and precise volumetric soft-tissue resection. RATS integrates OCT-guided 3D imaging with a calibrated fiber-coupled surgical laser and a sampling-based model predictive control (MPC) planner for constraint-aware resection. Our key contributions are:

- **Surgical Platform:** A modular RATS system integrating imaging (OCT and RGB-D) with an optically coupled co-axial fiber-coupled laser scalpel, calibrated via a multi-stage pipeline to align OCT, laser, and end-effector frames with sub-millimeter accuracy.
- **LTI modeling:** A data-driven, OCT-guided LTI model validated on tissue phantoms, using a super-Gaussian laser beam formulation to capture ablation crater morphology and tissue response.
- **Volumetric Resection Planning:** A sampling-based MPC framework operating directly on OCT voxel data, capable of generating constraint-aware resection trajectories and executing closed-loop ablations.

Several previous efforts have addressed isolated aspects of this pipeline, but have remained within the confines of simulation, single-point, or line-based resection. Ross et al. [9] proposed a genetic algorithm for raster-based surface ablation, though their method was limited to 2D surface evaluation and lacked sub-surface awareness. Wang et al. [10] proposed MPC-based planning; their system was validated only in simulation and did not integrate real-time sensing. Acemoglu et al. [12], [17] introduced data-driven LTI models for single-point or line-based ablations but lacked full volumetric integration. Prakash et al. [15] demonstrated a dual-sensor robotic platform for large-area OCT 3D volume reconstruction, but without integrated planning or resection.

To our knowledge, this is the first system to demonstrate calibrated, OCT-guided volumetric robotic laser resection in anatomically relevant settings. RATS tightly integrates perception, planning, and actuation for large-area (centimeter-scale) subsurface-aware (millimeter-scale) laser surgery, advancing the frontier of autonomous surgical robotics with potential applications in neuro-oncology, head and neck tumor removal, and soft tissue sarcoma resection requiring high precision and obstacle avoidance.

II. METHODS

A. System Overview

The RATS platform is a modular opto-mechanical robotic system designed for high-resolution imaging and precise, non-

TABLE I
MEAN AND STANDARD DEVIATION OF MEASURED LASER ENERGY.

Duty Cycle (%)	20	30	40	50	60	70	80	90	100
Mean Energy (J)	1.91	3.50	5.11	6.75	7.68	8.29	9.31	10.03	11.05
Std. Dev. ($\times 10^{-2}$ J)	1.6	1.4	3.3	3.0	6.4	3.4	12.0	10.3	23.7

contact laser-based tissue resection. The system integrates:

- **Imaging stack:** A low-cost spectral-domain (SD) OCT system, 1310 nm, (Lumedica Inc, NC, USA), with a working distance of 110 mm, is co-axially aligned with a near-infrared fiber-coupled surgical laser-scalpel using a dichroic mirror for synchronized imaging and ablation. A secondary RGB-D camera (Intel RealSense D435i) is used for coarse scene reconstruction and ROI selection. All components are mounted to a 6-DoF collaborative robotic arm (UR5e, Universal Robots, Denmark), which provides repeatable pose control for multi-angle scanning, laser delivery, and end-effector motion during volumetric resections.
- **Laser optics and scalpel assembly:** A surgical laser (λ =1060 nm, 10 W, multi-mode continuous wave, Crystal laser, NV, USA) is collimated (5 ± 1 mm) and focused using an achromatic doublet focusing lens (f = 75 mm, AC254-075-C Ø1", Thorlabs, NJ, USA). A long-pass dichroic mirror with an edge wavelength of 1072.4 nm (LPD02-1064RU-25x36x2.0, IDEX Semrock, IL, USA) co-aligns the OCT and laser scalpel paths, and a $\lambda/4$ UV fused silica Anti-Reflection (AR) coated protective window (TECHSPEC, Edmund Optics, NJ, USA) is installed at the exit port of the dichroic cage to protect the optics from debris generated during ablation. Paper burn tests are employed to measure the spot size at the focal length to be 0.9 mm. Total energy delivered over a set period of time under various duty cycles is measured using a power meter Fig. 4.
- **Phantom development:** The laser scalpel and OCT have different wavelengths, so a tissue-mimicking phantom was developed using 2% (w/w) agarose for matrix, 10% (w/w) intralipid for OCT optical scattering, and 0.1% (w/w) India-Ink as a chromophore to absorb 1060 nm laser energy. This phantom simulates the optical properties of soft tissue and is OCT-visible.
- **Controls architecture:** A Python-based control architecture manages robotic positioning, OCT scanning, laser trigger control, and closed-loop feedback. The interface supports pre-programmed patterns, trajectories, and voxel-level data from OCT, which are used to compute tissue resection paths in simulation and execution.

B. Multi-Sensor Calibration

Accurate volumetric resection requires precise alignment between the robot, imaging (OCT) and laser-ablation modules. We implement a multi-stage calibration procedure to estimate the rigid transformations necessary for closed-loop control and voxel-level accurate targeting.

1) Camera to End-effector Calibration: We estimate the extrinsic transformation between the RGB-D camera and the robot end-effector, T_{camera}^{EE}, using the OpenCV hand–eye calibration toolbox. A standard checkerboard is rigidly mounted
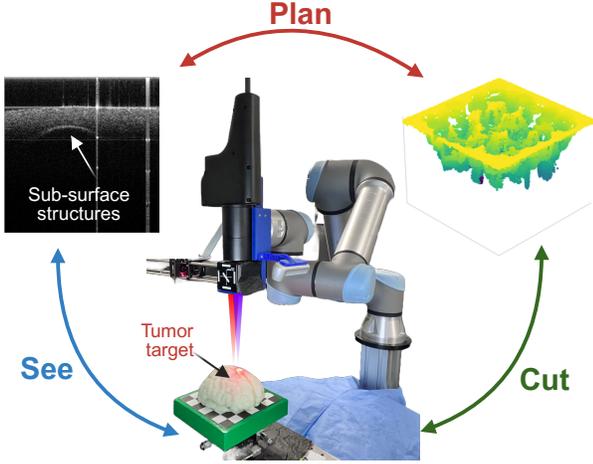
Fig. 2. **System Overview:** Experimental setup with a brain phantom for demonstration. The OCT sensor *sees* the surface and sub-surface structures, and the reconstructed 3D structure is then used to generate a volumetric resection *plan*. The planner generates a sequence of robot states and laser parameters to *cut* the tissue with clinically relevant objectives. [18]
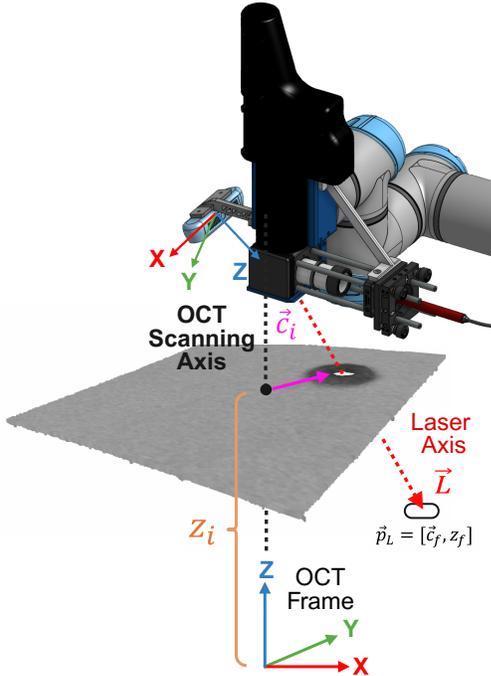

Fig. 3. **Laser Calibration Method:** Calibration schematic showing OCT frame (blue–green–red), OCT scanning axis, and laser axis alignment. The intersection point defines the laser focal point $\vec{p}_L$, estimated from crater centers $\vec{c}_i$ at different depths z_i.

in the workspace and observed from multiple robot poses. With the camera intrinsics pre-calibrated, the $AX = XB$ formulation [19] is used to solve for $X = T_{\text{camera}}^{EE}$, where A corresponds to robot motion and B to camera-observed checkerboard motion. This calibration enables coarse visual guidance and coordinate fusion.

2) OCT to End-Effector Calibration: Before 3D OCT volumes (C-scans) can be interpreted in world coordinates, the pixel-to-millimeter scale must be determined. To estimate lateral and axial scaling factors, we scan fiducial markers of known dimensions placed at the OCT focal plane. To

register the OCT frame to the robot end-effector, we estimate the transformation T_{OCT}^{EE} using the $AX = XB$ formulation [19], [20]. For any robot orientation, i, the matrix, $^{(i)}T_{EE}^{\text{world}}T_{OCT}^{EE}\,^{(i)}T_{\text{marker}}^{OCT}$, is constant as it relates a point in a stationary frame relative to the calibration pattern to its coordinates in the world frame. Hence, obtaining multiple values of T_{EE}^{world} and T_{marker}^{OCT} and rearranging them into A and B matrices provides an estimate of $X = T_{OCT}^{EE}$ [19].

While the T_{EE}^{world} matrix is easily obtained from robot joint angles and forward kinematics, Our method for estimating the T_{marker}^{OCT} matrix uses a two-dot calibration board paralleling [20] (Fig. 6). We define the stationary "marker" frame by denoting the vector between the two dots as the $+X$ axis, while the vector orthogonal to the marker plane is denoted as the $+Z$ axis. The $+Y$ axis can then be obtained from a cross product operation between the two vectors. From knowing these axes in the OCT frame, the OCT to marker transformation can be obtained (T_{marker}^{OCT}).

3) Laser to OCT Extrinsic Calibration: To track the laser, we calibrate directly to the OCT system by estimating T_{laser}^{OCT}. This matrix can be directly estimated by obtaining the vector $\vec{L}$ that points along the laser axis in the OCT frame, and the location of the laser focal point, $\vec{p}_L$. This can be obtained by performing a sequence of ablations at different heights, z_i, along the OCT axis, then scanning using OCT and estimating the center $\vec{c}_i = [x_i, y_i]$ of each ablation crater on the XY plane in the OCT frame. These points form a parametric line in $\mathbb{R}^3$ representing the laser axis, $\vec{L}$, which can be obtained from calculating a line-of-best-fit through collected $\vec{c}_i$ and z_i pairs. To estimate $\vec{p}_L$, we can find the point along the line $\vec{L}$ corresponding to a known focal distance, z_f (Fig. 3).

C. Laser Tissue Interaction Model

To generate a tissue ablation plan, the tissue's response to varied geometry and energy distribution of laser-induced ablation craters is needed. We model single-point tissue resection using a two-dimensional Gaussian beam model [9], [10] and a steady-state ablation formulation [4]. This enables us to estimate interpretable tissue parameters, such as tissue density and enthalpy, as a function of incident energy, resulting in the prediction of the ablated tissue crater.

Gaussian profile estimation and interpretation: The 3D ablation crater can be represented geometrically by conditioning the steady state laser ablation model upon the laser beam position. With the incident laser beam centered at (μ_x, μ_y) on the tissue surface, the resulting ablation crater depth, f, at a point, (x, y), on the surface can be calculated as

$$f(x, y) = -A\max\left(0, E\exp\left[-\left(\frac{r^2}{2\sigma^2}\right)^P\right] - \phi\right), \quad (1)$$

$$r^2 = (x - \mu_x)^2 + (y - \mu_y)^2, \quad (2)$$

where A is an amplitude dependent on laser/tissue parameters, σ is the standard deviation of the distribution (the radius where the intensity falls to approximately 60.6% of the maximum value), ϕ is an offset parameter representing the minimum energy required for cutting to occur, E is the total energy
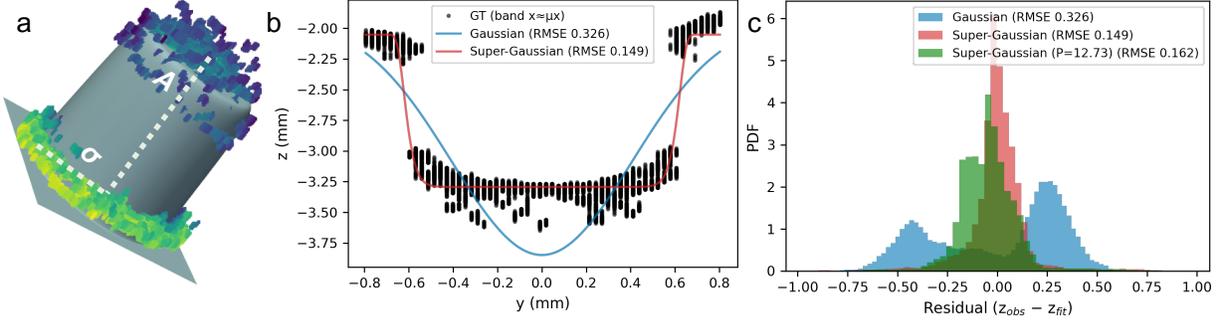
Fig. 4. **Laser–tissue interaction (LTI) model estimation. (a)** Example of super-Gaussian parameter estimation through fitting a curve to the ground-truth point cloud obtained from OCT B-scans. This process enables estimation of single-point tissue response for a given laser energy. **(b)** Example of Gaussian and super-Gaussian curve fits on a ground-truth point cloud along the x-axis. Gaussian curve estimation does not capture the multimode beam response adequately. **(c)** Probability density function of residual error from curve fitting (z_{fit}) with respect to observed (z_{obs}) point cloud for Gaussian and super-Gaussian fits, at 50% duty cycle with the LTI model ($P = 12.73$).

delivered, and $P \in \mathbb{R}$ modulates the profile sharpness. When $P = 1$, the profile reduces to a standard Gaussian, which was used in [9], [10], [21], but was insufficient for our setup. For $P > 1$, the crater has a flatter bottom and steeper sidewalls, which is often observed in high-power, quasi-top-hat beams.

D. Parameter regression

Parameters are estimated via the Levenberg–Marquardt algorithm [22], which adaptively interpolates between gradient descent and Gauss–Newton updates for efficient convergence in nonlinear parameter spaces. To reduce the influence of edge noise and overfitting, we only fit to a circular region within 2σ of the crater center. Given cutting errors, $e_i = z_i - f(x_i, y_i)$, the model's accuracy is evaluated using the root mean square error (RMSE), $N^{-0.5}(\sum_{i=1}^{N} e_i^2)^{-0.5}$, percent overcut (%OC), $\sum_{i=1}^{N} \min(0, e_i) / \sum_{i=1}^{N} e_{i,0}$, undercut (%UC), $\sum_{i=1}^{N} \max(0, e_i) / \sum_{i=1}^{N} e_{i,0}$, and intersection over union (IoU),

$$\text{IoU} = \frac{V_{real} \cap V_{fit}}{V_{real} \cup V_{fit}} = \frac{\sum_{i=1}^{N} \max(z_i, f(x_i, y_i))}{\sum_{i=1}^{N} \min(z_i, f(x_i, y_i))},$$

where z_i, $f(x_i, y_i)$, and $z_{i,0}$ denote the ablation depth, target depth, and initial depth at x_i, y_i, and $e_{i,0} = z_{i,0} - f(x_i, y_i)$.

Special importance is given to % overcut, as ablated tissue cannot be "uncut", meaning overcut leads to difficulty and unremovable error in future planning.

E. Volumetric Resection Path Planning Algorithm

Previous attempts at laser-based volumetric tissue resection have primarily relied on open-loop raster patterns or fixed-depth slicing, limiting success when used with subsurface features or anatomical constraints. In contrast, we present the first experimental implementation of a closed-loop, constraint-aware, intelligent volumetric resection planner capable of executing volumetric resection in both feedforward and feedback modes. The implementation presented in this work accounts for several practical challenges of non-contact, laser-based surgery, such as the inevitable presence of debris, geometric workspace constraints, imaging system artifacts, and estimation of performance error. OCT is used as the

primary feedback sensor due to its micron-scale resolution and robustness in soft tissue imaging. OCT volume workspace is discretized, including the target (tumor) and hard constraints (critical structures to avoid). Clough-Toucher splines are used to interpolate between points and generate smooth surfaces.

The planner employs the sampling algorithm from [10] to select ablation inputs based on a single-shot laser–tissue interaction model. The algorithm samples the input space of possible cuts after being provided with an objective surface (tumor to remove) and a constraint surface (critical regions to avoid), which can be obtained either from OCT scanning or preoperative imaging. It constructs a tree with nodes containing predicted tissue surfaces and associated costs

$$C(\mathcal{X}) = \sqrt{\frac{1}{N} \sum_{i=1}^{N} \left(\min(0, e_i) + \lambda \max(0, e_i) \right)^2} \quad (3)$$

with weighting factor $\lambda \geq 1$ to favor undercuts, and edges containing candidate inputs. The candidate inputs are parameterized as $\vec{u} = (\vec{x}_L, \vec{\theta}_L, D)$, where $\vec{x}_L = (\mu_x, \mu_y)$ defines the laser incident point on the tissue surface, $\vec{\theta}_L = (\theta_x, \theta_y)$ specifies the laser incident angle, and D is the laser duty cycle, assuming a constant nominal laser power of 8W and the ablation duration is held constant at 1.5s. Initially, the tree contains only the original tissue surface. Subsequent edges are added to nodes selected randomly, but weighted towards lower costs.

Each edge's input is selected at random from the set of allowable laser positions, angles, and duty cycles. The resulting child node contains the simulated tissue surface generated by applying the selected laser cut to its parent node's surface. If the resulting state is found to violate any constraints, the node/edge are not added. This procedure repeats for a set number of iterations (k_F), before an input sequence is chosen as the path between the root node and the cost node in the tree. In feedforward mode, this algorithm is repeated, with the prior end-node as the root node of a new tree. The algorithm is also repeated in feedback mode, but after performing a sequence of m cuts, a new OCT scan provides an observed surface for the root node of a new tree.
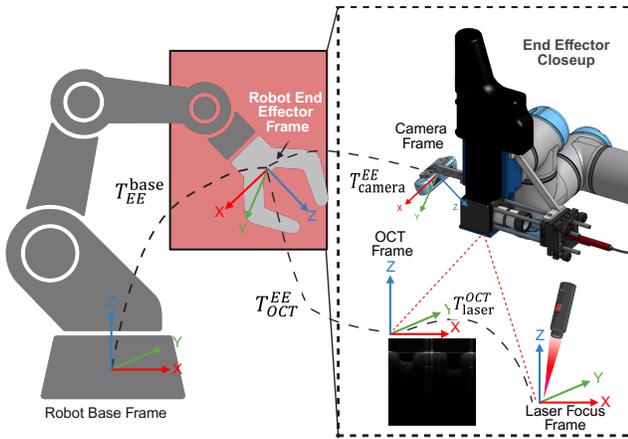
Fig. 5. **System calibration pipeline and coordinate frame hierarchy:** The robot end-effector houses the OCT, RGB-D camera, and laser scalpel. Transformations between robot base, end-effector, OCT, camera, and laser focus frames are shown. Includes OCT 2D-to-3D volume registration, OCT-to-end-effector (EE) frame transformation, and laser-to-OCT alignment. Dashed lines indicate calibrated extrinsic transforms (T_{OCT}^{EE}, T_{laser}^{OCT}, T_{camera}^{EE}). The figure illustrates the reference frame alignment critical for precise imaging and laser actuation.

III. EXPERIMENTS AND RESULTS

We evaluated the RATS platform, LTI model, and MPC-based planner to assess accuracy, robustness, and surgical applicability. Experiments included (i) system calibration, (ii) LTI model fitting, and (iii) volumetric resection on tissue phantoms and *ex vivo* porcine samples with and without feedback.

A. System Verification and Error analysis

1) OCT calibration : Using fiducial markers, the OCT pixel-to-millimeter spacing was determined to be 14 μm laterally and 14.59 μm axially. For each C-scan, 256 B-scans were collected with a distance of 28 μm between adjacent B-scans [15], [20]. Each B-scan consisted of 512 A-scans, and the resulting images were processed to obtain a 3D tissue surface using the methods described in [12], [15]. Each C-scan encompassed a $7.168 \times 7.168 \times 7.6288$ mm³ volume, with axial penetration limited to 2–4 mm depending on the tissue's optical properties. Fig. 6 summarizes the calibration error analysis for OCT-to-EE and laser-to-OCT alignment. The mean reprojection errors for points used in the OCT-to-EE calibration were 0.638 mm for the origin and 0.625 mm for the x-axis, with occasional outliers above 2 mm. Fig. 6a .

2) Laser Calibration Verification: A 3×3 dot grid target was scanned by the OCT to generate a C-scan and a rasterized en-face image, enhancing dark circular blobs. Blobs were segmented using standard image processing techniques, including thresholding and morphological operations. The centroids of each blob were extracted and then mapped back to the metric OCT frame. The z-coordinate at each center was obtained from a global plane fit to the full surface point cloud.

For laser scalpel calibration verification, four targets were selected, and the laser was used to resect the targets for 0.5 s while maintaining orientation along the surface normal. The remaining targets were used to evaluate OCT repositioning error across successive scans (Fig. 6b). The measured laser

calibration error was 0.161 ± 0.031 mm, well below the 0.5 mm neurosurgical threshold, confirming calibration accuracy. It is worth noting that the OCT repositioning error to a previously scanned point was 0.071 ± 0.009 mm, further validating the calibration pipeline. The larger ablated circles observed in Fig. 6b arose from selectively higher absorption of the 1060 nm laser by black-dot chromophores, which enhanced local heat dissipation.

B. LTI model performance

Accurate estimation of LTI parameters is critical for downstream ablation planning. Agar phantoms were prepared according to Section II-A and cured for four hours prior to use. An OCT C-scan was then acquired to measure surface inclination with respect to the robot arm. The laser scalpel was aligned with the tissue surface normal, and point ablations were performed at a minimum of three locations for each duty cycle (Fig. 4a). The ablation duration was fixed at 1.5 s across all trials. To ensure consistent and high-quality data, the protective window was cleaned between ablations to remove any debris that may have been accumulated. Debris generation is a well-documented phenomenon in laser tissue ablation, arising from microbubble explosions [23]; debris on the protective window attenuates and diffuses laser energy, resulting in inconsistent craters if not removed.

Following the procedure described in Section II-C, LTI parameters were estimated and presented in Table II from experiments across multiple duty cycles. The super-Gaussian sharpness parameter P was first treated as a free variable and then fixed to its mean value for subsequent fitting. As expected, given the multimodal profile of the 1060 nm surgical laser (with a 650 nm red alignment beam), the super-Gaussian (SG) model consistently outperformed Gaussian fits in capturing crater morphology. Final parameters used for simulation are highlighted in blue, Table II. Independent phantom validation yielded an RMSE of 0.402 mm between measured and simulated data, confirming predictive accuracy. The ablation threshold parameter ϕ was determined experimentally as the lowest energy producing a cut (20% duty cycle, $\phi = 1.939$ J).

The SG model consistently achieved lower RMSE values than Gaussian fits, with the greatest improvements observed at higher duty cycles where the multimode laser profile produced flat-bottom craters and steep sidewalls. For instance, at 50% duty cycle, the SG model achieved RMSE values of 0.169–0.242 mm compared to 0.326–0.386 mm for Gaussian fits. On average, the SG model yielded a sharpness parameter of $P = 12.73$, amplitude $A = 0.334 \pm 0.048$, and beam spread $\sigma = 0.473 \pm 0.054$ mm, consistent with quasi-top-hat beam characteristics. These findings demonstrate that the SG formulation provides a more accurate representation of multimodal laser–tissue interactions, enabling improved prediction and planning for volumetric resections.

Fig. 4 further illustrates this comparison. In Fig. 4a, crater morphology is reconstructed from OCT point clouds, with amplitude A and beam spread σ fitted to describe ablation depth and lateral spread. Fig. 4b compares Gaussian and SG fits along a crater cross-section: the Gaussian model
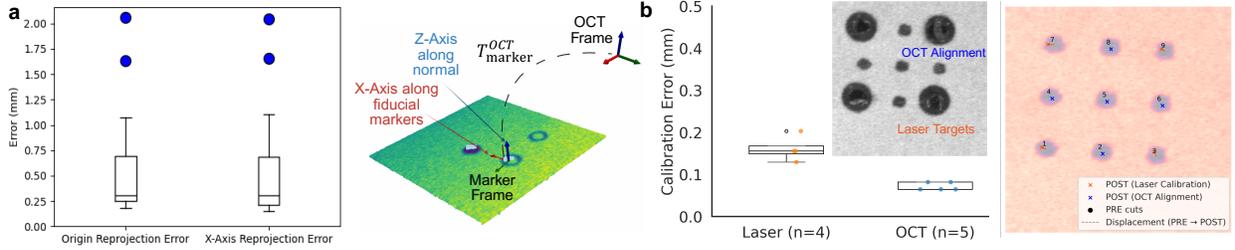
Fig. 6. **Calibration Error Analysis: (a)** OCT-to-end-effector calibration accuracy (left) and relevant coordinate frames (right). **(b)** Laser-to-OCT calibration and representative *en-face* image after laser calibration test (left), OCT-probe repositioning verification error (right).

TABLE II
FITTING RESULTS FOR GAUSSIAN VS. SUPER-GAUSSIAN (SG) MODELS

PWM	Trial	P (SG)	A	σ	RMSE (SG)	RMSE (G)
30	1	3.38	0.428	0.390	0.141	**0.105**
	2	2.67	0.344	0.431	0.119	**0.078**
	3	2.98	0.341	0.379	**0.104**	0.122
50	1	22.11	0.357	0.526	**0.242**	0.386
	2	20.74	0.333	0.519	**0.169**	0.357
	3	10.58	0.390	0.489	**0.199**	0.326
70	1	8.97	0.270	0.524	**0.242**	0.405
	2	16.63	0.252	0.517	**0.169**	0.375
	3	24.06	0.295	0.487	**0.199**	0.465
90	1	17.26	0.322	0.539	**0.281**	0.680
	2	10.84	0.324	0.505	**0.373**	0.666
	3	12.52	0.350	0.492	**0.534**	0.812
Average		12.73	0.334	0.473	0.231	0.398
Std. Dev.		7.52	0.048	0.054	0.121	0.223

TABLE III
ABLATION RESULTS COMPARING FEEDFORWARD (NO FEEDBACK) AND
CLOSED-LOOP (FEEDBACK) PLANNING IN SIMULATION AND REAL
EXPERIMENTS.

Condition		Obj. Vol. (mm^3)	RMSE (mm)	MAE (mm)	%OC	%UC	IoU
No Feedback	Simulation	71.4	0.798	0.554	15.0	24.7	62.3
	Real	71.4	1.007	0.797	9.8	47.3	21.1
Feedback	Simulation	71.4	1.199	0.842	37.1	23.3	60.0
	Real	71.4	0.842	0.630	19.3	25.8	60.0

underestimates the flat-bottom region, while the SG model closely matches the ground truth, yielding a lower RMSE (0.149 vs. 0.326). Finally, Fig. 4c shows residual distributions, where the SG fit produces a narrower, more symmetric profile even with average model parameters, confirming improved predictive accuracy.

C. Volumetric Resection through MPC-Planner

The surgical planner described in Section II-E was evaluated across three increasingly complex tasks. First, it was used to resect phantom tissue with user-defined objectives and constraints to simulate a challenging resection scenario. Second, the constraint was defined by the presence of a subsurface critical structure extracted from an OCT C-scan. Third, the planner was deployed on *ex vivo* porcine tissue. Across all experiments, the permissible angular tilt of the laser was limited to $10°$ to prevent collisions within the workspace.

1) Phantom Resection: Feedforward mode: A square well of $6 \times 6 \times 2$ mm^3 was defined in simulation, and an ablation plan was generated. The dimensions were chosen to approximate the typical size of brain tumor samples, where other laser-based thermal therapies, such as laser interstitial thermal therapy, are not applicable. The resection plan was terminated once the residual tumor volume fell below 25% of

the objective. The same plan was then executed on the RATS system using a tissue phantom. Because it was impractical to clean the protective window after every ablation step, the window was inspected for debris every three steps. Upon completion, a post-ablation OCT scan was acquired. The simulated plan, post-ablation volume, and associated errors are presented in Fig. 7 and Table III. The RMSE and MAE errors increased by 26.19% and 43.86% over simulated results when executed.

2) Phantom Resection: Feedback mode: Using the same square, well-shaped target, an ablation plan was generated and re-planned every 9 steps to compensate for thermal and debris effects. The final performance metrics are presented in Table III. Compared to feedforward mode, feedback improved IoU agreement between the simulation and experiments by 64.8%.

In the real no-feedback case, a significant undercut was observed, likely caused by attenuation of the laser beam due to loss of focus as the cuts deepened. This effect was mitigated when feedback was incorporated, demonstrating the effectiveness of closed-loop control. The RMSE and MAE error of 1.007 mm and 0.797 mm without feedback and 0.842 mm and 0.630 mm with feedback, confirming that both methods achieves clinically acceptable accuracy within the 1 mm neurosurgical threshold. Unlike the previous case, when feedback was included, both RMSE and MAE in the real experiment improved (decreased) by 29.77% and 25.17% over the simulation. It should be noted that RMSE values may be inflated by OCT speckle noise.

3) Resection with a Subsurface Critical Structure: Neurosurgical tumors often envelop critical structures, such as blood vessels, and thus every unnecessary resection step potentially deprives the patient of key bodily functions. Thus, identification and preservation of critical anatomical structures is a *must* for a successful resection plan.

To mimic a subsurface structure visible in OCT, a ball bearing (3.175 mm OD) was embedded beneath layers of the phantom, chosen for its optical distinctiveness. The contour of the bearing was segmented from the OCT B-scans alongside the tissue surface using the DBSCAN algorithm [24] for clustering, allowing the user to label the critical structure. The segmented subsurface structure was isolated and designated as a hard constraint that needed to be avoided during resection planning. The phantom surface was uneven, resembling tissue-like variability (Fig. 8a, b). Unlike the previous square-well experiment, here an artificial tumor shape was defined
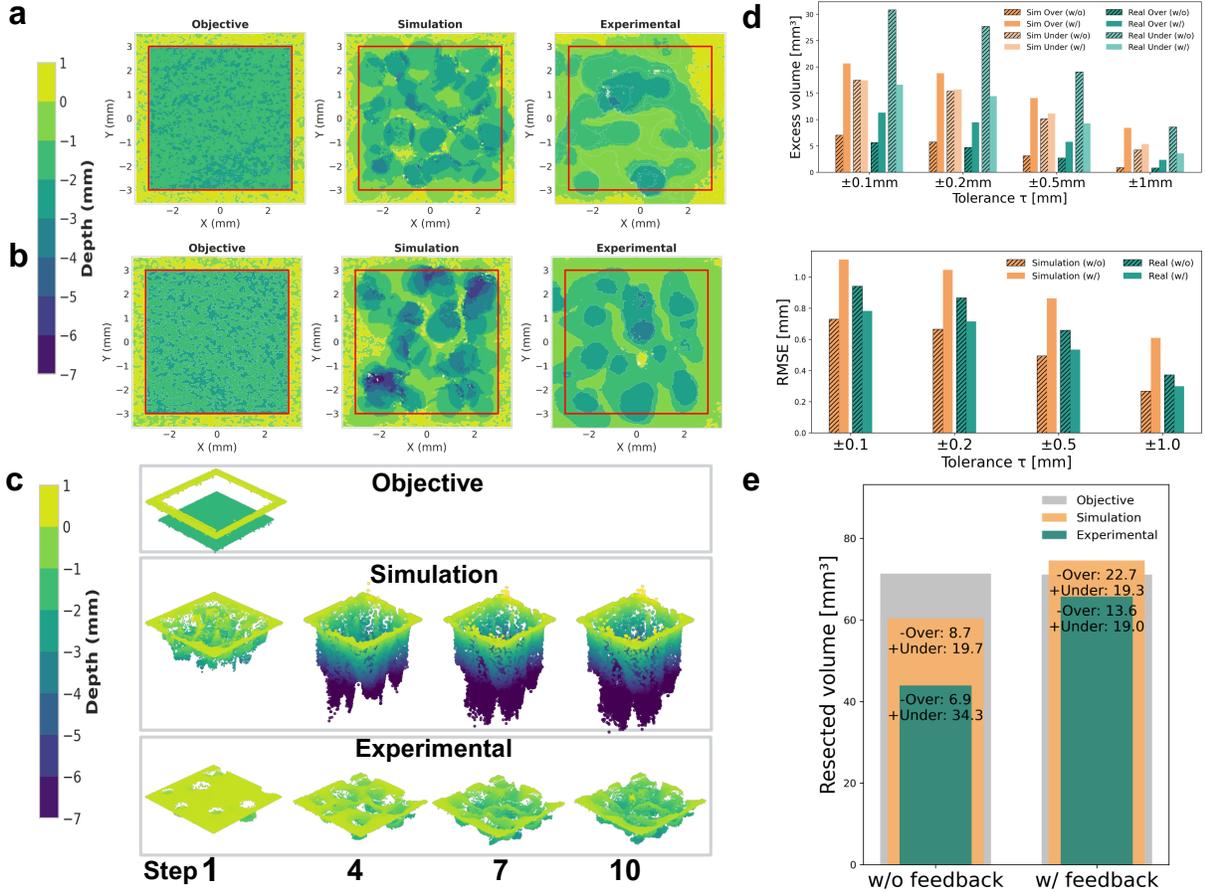
Fig. 7. **Volumetric resection planner (a)** without feedback , simulation and experimental result in 2.5D with depth information at each XY point representing color. The red bounding box denotes the objective boundary. The experimental result can be seen undercutting; **(b)** with repeated OCT feedback. Simulation shows considerable overcut, whereas the experimental result is closer to the objective; **(c)** intermediate steps of (b) showcasing the process of incorporating regular feedback. Experimental resection progresses in controlled manner; **(d)** planner performance in terms of excess volume (top) and RMSE (bottom), over clinical directional tolerance; **(e)** resected volume in comparison to the objective in the simulation and experimental case.

using an irregular closed spline in the XY-plane centered at the tumor location, with mean radius 1.5 mm and random radial jitter to produce a non-circular footprint. The maximum permissible resection depth was limited to $2\,\mathrm{mm}$. A continuous subsurface constraint volume was developed by interpolation using Clough-Toucher algorithm, and an clearance margin of 0.2 mm was added to the constraint to enforce safety. The objective function was generated following the constraints, and a resection plan was generated with feedback and executed. As seen in Fig. 8a, the embedded subsurface structure was within $0.5mm$ of the top surface, resulting in the planner making minor, less perceivable cuts.

4) Ex vivo tissue resection: Finally, the system was evaluated on an *ex vivo* porcine abdominal tissue, Fig. 8c, d. The surface was coated with an absorbing chromophore (India Ink) to lower the ablation threshold to 1060 nm and induce initial ablation. Using phantom-derived LTI parameters without retraining, RATS achieved a $4 \times 4 \times 2$ mm^3 resection. While accuracy was reduced compared to phantom trials due to debris, scattering, and a lack of a porcine-specific LTI model, the experiments demonstrated feasibility and offered critical insights into the difficulties of working with complex porcine tissue. These findings highlight robustness under anatomically realistic conditions and motivate the development of tissue-

specific LTI models. Note that the above performance can be improved at the cost of additional time and a tissue-specific LTI model.

IV. Discussion And Conclusion

This work introduces RATS, the first closed-loop robotic platform for OCT-guided volumetric laser resection. By integrating OCT and RGB-D imaging, sub-millimeter multi-stage calibration, and a sampling-based MPC planner with closed-loop feedback, RATS achieved clinically relevant performance. The system demonstrated calibration accuracy of 0.161 ± 0.031 mm, robust LTI modeling (RMSE 0.231 ± 0.121 mm), and volumetric resection accuracy of 0.842 mm RMSE with feedback. IoU agreement improved by 64.8% compared to feedforward execution, and constraint-aware planning reliably preserved subsurface structures in phantom and *ex vivo* experiments. These results validate the feasibility of autonomous, obstacle-aware volumetric laser resections.

Our results highlight the possibility of obtaining clinically relevant performance with practical hardware design, an empirical LTI model, and a modular sampling-based MPC planner for volumetric resection. The proposed calibration methods extend beyond the proposed application to other free-form or fiber-coupled laser systems. The tissue phantom
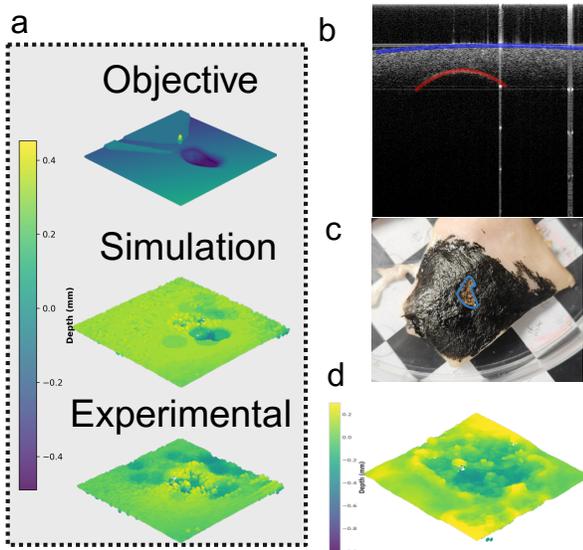
Fig. 8. **Subsurface Resection and Tissue: (a)** Resection objective, simulation, and experimental results for phantom with subsurface critical structure. The objective function is defined by taking the constraint into account. As the constraint is close to surface, shallow cuts can be observed; (b) OCT B-scan showing surface (blue) and constraint (red); **c** porcine tissue sample post-resection, stained with India Ink dye; **d** resected result from porcine tissue.

developed is from easy-to-source, biocompatible materials to enhance the laser scalpel's performance while accounting for suitability with OCT-based scanning. The entire codebase is implemented in Python on an Intel i9 CPU (11th gen, Intel Corporation, CA, USA) and has the potential to be further sped up with a C-based implementation in the future.

Despite these advances, several limitations remain. First, debris accumulation during ablation reduces laser transmission and accuracy, motivating the integration of pressurized air flushing and extended working distances. Second, the current laser spot size (0.9 mm) limits boundary precision; improved optics targeting < 0.5 mm are required for fine resections. Third, the present LTI model is geometric and does not account for thermal accumulation across repeated ablations, which can alter tissue properties. Finally, the sampling-based planner is slower than raster strategies, though it enables constraint-aware resections in geometrically complex cases.

Clinically, RATS highlights the potential of robotic laser systems for precise, safe, and constraint-aware resections in neurosurgery, head and neck oncology, and soft-tissue resection. This work advances the frontier of autonomous surgical robotics and lays the foundation for translation toward in vivo tumor resections.

In conclusion, we presented RATS, a novel robotic platform integrating OCT imaging and laser-based resection for intelligent, precise, and safe surgical intervention. The proposed work is the culmination of several independent challenges, ranging from hardware, modeling, and testing, and is an effort to show the final surgical performance upon implementation for volumetric resection tasks. Future work will undertake improvements in opto-mechanical design for surgical deployment, a generalizable LTI model [12], and real-time adaptive planning for achieving sub-mm surgical performance.

REFERENCES

[1] G. P. Moustris, S. C. Hiridis, K. M. Deliparaschos, and K. M. Konstantinidis, "Evolution of autonomous and semi-autonomous robotic surgical systems: a review of the literature," *The international journal of medical robotics and computer assisted surgery*, vol. 7, no. 4, pp. 375–392, 2011.

[2] P. Fiorini, K. Y. Goldberg, Y. Liu, and R. H. Taylor, "Concepts and trends in autonomy for robot-assisted surgery," *Proceedings of the IEEE*, vol. 110, no. 7, pp. 993–1011, 2022.

[3] H. C. Lee, N. E. Pacheco, L. Fichera, and S. Russo, "When the end effector is a laser: A review of robotics in laser surgery," *Advanced Intelligent Systems*, vol. 4, no. 10, p. 2200130, 2022.

[4] A. Vogel and V. Venugopalan, "Mechanisms of pulsed laser ablation of biological tissues," *Chemical reviews*, vol. 103, no. 2, pp. 577–644, 2003.

[5] H. C. Lee, N. E. Pacheco, L. Fichera, and S. Russo, "When the end effector is a laser: A review of robotics in laser surgery," *Advanced Intelligent Systems*, vol. 4, no. 10, 2022.

[6] P. A. York, R. Peña, D. Kent, and R. J. Wood, "Microrobotic laser steering for minimally invasive surgery," *Science Robotics*, vol. 6, no. 50, p. eabd5476, 2021.

[7] A. Gunalan and L. S. Mattos, "Towards oct-guided endoscopic laser surgery—a review," *Diagnostics*, vol. 13, no. 4, p. 677, 2023.

[8] S. Li, A. Gunalan, M. A. Azam, V. Penza, D. G. Caldwell, and L. S. Mattos, "Auto-calm: Autonomous computer-assisted laser microsurgery," *IEEE Transactions on Medical Robotics and Bionics*, 2024.

[9] W. Ross, N. Cornwell, M. Tucker, B. Mann, and P. Codd, "Optimized path planning for soft tissue resection via laser vaporization," in *Clinical and Translational Neurophotonics 2018*, vol. 10480. SPIE, 2018, pp. 16–22.

[10] V. Y. Wang, R. Prakash, S. R. Oca, E. J. LoCicero, P. J. Codd, and L. J. Bridgeman, "Sampling-based model predictive control for volumetric ablation in robotic laser surgery," *arXiv e-prints*, pp. arXiv–2410, 2024.

[11] R. Prakash, "Created in biorender," https://BioRender.com/qhfnios, n.d., accessed: 2025-09-15.

[12] G. Ma, R. Prakash, B. Mann, W. Ross, and P. Codd, "3d laser-and-tissue agnostic data-driven method for robotic laser surgical planning," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 8446–8453.

[13] Q. Lan, H. Zhang, W. Zhang, C. Gu, and J. Yang, "Oct as both a shape sensor and a tomographic imager for large-scale freeform robotic scanning," *Opt. Lett.*, vol. 50, no. 1, pp. 45–48, Jan 2025. [Online]. Available: https://opg.optica.org/ol/abstract.cfm?URI=ol-50-1-45

[14] D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, C. A. Puliafito, *et al.*, "Optical coherence tomography," *science*, vol. 254, no. 5035, pp. 1178–1181, 1991.

[15] R. Prakash, V. Wang, Z. Chen, L. Bridgeman, and P. J. Codd, "Portable dual sensor large area visualization system for robotic laser surgery," in *2025 International Symposium on Medical Robotics (ISMR)*. IEEE, 2025, pp. 58–65.

[16] G. Ma, M. McCloud, Y. Tian, A. Narawane, H. Shi, R. Trout, R. P. McNabb, A. N. Kuo, and M. Draelos, "Robotics and optical coherence tomography: current works and future perspectives," *Biomedical Optics Express*, vol. 16, no. 2, pp. 578–602, 2025.

[17] A. Acemoglu, L. Fichera, I. E. Kepiro, D. G. Caldwell, and L. S. Mattos, "Laser incision depth control in robot-assisted soft tissue microsurgery," *Journal of Medical Robotics Research*, vol. 2, no. 03, p. 1740006, 2017.

[18] R. Prakash, "Created in biorender," https://BioRender.com/lz5npof, n.d., accessed: 2025-09-15.

[19] R. Tsai, "A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses," *IEEE Journal on Robotics and Automation*, vol. 3, no. 4, pp. 323–344, 2003.

[20] G. Ma and M. Draelos, "Geometry-aware volumetric data stitching using local surface mapping and robot optical coherence tomography," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 3788–3794.

[21] N. E. Pacheco, C. S. Gaddipati, S. Farzan, and L. Fichera, "Automatic focus adjustment for single-spot tissue temperature control in robotic laser surgery," *IEEE transactions on medical robotics and bionics*, 2024.

[22] H. P. Gavin, "The levenberg-marquardt algorithm for nonlinear least squares curve-fitting problems," *Department of Civil and Environmental Engineering Duke University August*, vol. 3, pp. 1–23, 2019.

[23] R. Verdaasdonk, C. Borst, and M. Van Gemert, "Explosive onset of continuous wave laser tissue ablation," *Physics in Medicine & Biology*, vol. 35, no. 8, p. 1129, 1990.

[24] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *kdd*, vol. 96, no. 34, 1996, pp. 226–231.