Automated and deterministic Trajectory Planning for Stereotactic Brain Biopsy Using Patient-Specific Anatomical Modeling and Exhaustive Discrete Search

1st Mohammad Jafari Vayeghan
School of ECE, College of Engineering
University of Tehran
Tehran, Iran
Mhmmd.jafari@ut.ac.ir

2nd Niloufar Delfan
Department of EECS
Lassonde School of Engineering
York University
Toronto, Canada
ndelfan@yorku.ca

3rd Sara Parvaresh Rizi
Faculty of Applied Science and Engineering
University of Toronto
Toronto, Canada
s.parvareshrizi@mail.utoronto.ca

4th Mansour Parvaresh Rizi
Department of Neurosurgery
School of Medicine
Iran University of Medical Sciences
Tehran, Iran
parvareshrizi.m@iums.ac.ir

5th Mehdi Tale Masouleh
School of ECE, College of Engineering
University of Tehran
Tehran, Iran
m.t.masouleh@ut.ac.ir

6nd Ebrahim Ghafar-zadeh
Department of EECS
Lassonde School of Engineering
York University
Toronto, Canada
egz@yorku.ca

7th Behzad Moshiri
School of ECE, College of Engineering
University of Tehran
Tehran, Iran
moshiri@ut.ac.ir

***Abstract*—Stereotactic brain biopsy for intracranial lesions carries inherent risks, with hemorrhagic complications occurring in 0.7–5.2% of cases, often due to suboptimal trajectory planning rather than intraoperative deviation. Manual planning is time-intensive (12–30 minutes) and operator-dependent, while existing automated methods rely on manual region-of-interest definition or stochastic optimization lacking global optimality guarantees. We developed a fully automated trajectory planning framework integrating: (1) deep learning-based segmentation of cerebral vasculature (NeuroVASC U-Net, DSC=86.09%) and tumor core (3D Attention U-Net, DSC=75.79%) from T1-weighted contrast-enhanced MRI; (2) geometric hairline constraint automatically defining anatomically valid skull entry zones (27,051 candidate points vs. 500–2,000 in manual ROIs); (3) deterministic exhaustive search guaranteeing global-optimal trajectories with multi-path reporting and multi-tumor support. Validation was performed on 137 consecutive patients. On a probe envelope of 5 mm the system achieved a 99.3% patient-level and 98.5% tumor-cluster-level planning success rate with a mean computation time of 10.54 seconds per patient, an approximately 10-fold reduction versus manual planning, while providing five spatially diverse alternative trajectories. At a 6 mm envelope, patient-level success was 92.7%, with explicit failure flagging in the remaining cases. Qualitative review by independent neurosurgeons confirmed clinical feasibility. This framework addresses critical limitations of existing methods through automated segmentation, intelligent search space pruning, and deterministic optimization, positioning it for prospective clinical validation.**

***Index Terms*—Biopsy systems & technologies, MRI-compatible instrumentation, AI-assisted image-guided interventions, Deep Learning, Path Planning**

I. Introduction

Stereotactic brain biopsy is a cornerstone of modern neuro-oncology, providing minimally invasive histopathological diagnosis for intracranial lesions located in deep-seated regions, eloquent cortex, or areas not amenable to open resection [1]. Although advances in neuroimaging, frameless navigation, and stereotactic guidance have improved procedural accuracy, stereotactic biopsy remains associated with a measurable risk of morbidity and mortality. Reported complication rates range from approximately 1% to 8.5%, with procedure-related mortality between 0.6% and 1.34%, and intracranial hemorrhage constituting the most frequent and clinically significant adverse event [2], [3]. Accumulating evidence indicates that the majority of iatrogenic hemorrhages occur not as a consequence of intraoperative deviation from the planned trajectory, but rather due to suboptimal preoperative trajectory selection that fails to adequately account for patient-specific vascular anatomy and critical structures [4]. This observation underscores the central role of trajectory planning as a primary determinant of procedural safety [5].

Designing a safe and effective biopsy trajectory is inherently a multi-objective optimization problem that requires simultaneous consideration of anatomical, geometric, and diagnostic constraints [6]. An optimal trajectory must maximize clearance from cerebral vasculature, ventricles, and eloquent brain regions while maintaining a short path length and ensuring representative tissue sampling from the lesion [5]. In routine clinical practice, this process is performed manually through visual inspection of volumetric imaging, most commonly contrast-enhanced T1-weighted (T1CE) MRI, by selecting an entry point on the skull and a target within the lesion, followed by assessment of the straight-line path connecting them. While intuitive, this approach is time-intensive, highly operator-dependent, and subject to inter-surgeon variability.

To address the limitations of manual trajectory selection, automated path-planning frameworks have been developed for minimally invasive stereotactic neurosurgery [7], [8]. Despite implementation differences, most systems follow a common image-driven pipeline (Fig. 1), beginning with volumetric neuroimaging acquisition; typically T1CE MRI; supplemented by vascular, diffusion, or functional imaging [8].

Despite advances, existing automated methods still have limitations. Many require the surgeon to specify a candidate entry region or target point in advance. Some rely on stochastic search or metaheuristics, which do not guarantee finding the global optimum. The handling of safety constraints can be inconsistent (weighting vessels versus functional areas unevenly). Finally, most published planners have only been tested on relatively small or homogeneous patient cohorts, leaving their general clinical performance uncertain.

This study presents a fully automated, deterministic framework for stereotactic biopsy trajectory planning that combines patient-specific anatomical modeling with global optimization. Cerebral vasculature, tumor core, and CSF spaces are segmented from T1CE MRI using deep learning (DL) models, NeuroVASC U-Net [9] for vascular segmentation and a 3D Attention U-Net [10] for tumor delineation, augmented by semi-automatic methods. A novel geometric strategy automatically defines a broad, anatomically valid skull entry zone using skull surface constraints and hairline landmarks, eliminating manual ROI selection. Within this space, an exhaustive discrete search evaluates all feasible straight-line trajectories using Euclidean distance transform–based safety metrics, guaranteeing global optimality in the discretized domain. The framework additionally reports multiple spatially distinct alternative trajectories and supports sequential planning for multi-lesion cases through dynamic obstacle updates.

II. Materials and Methods

A. Data Acquisition and Annotation

This study analyzed a prospectively collected dataset of 1.5-T T1-T1CE MRI scans from 137 patients who underwent stereotactic brain biopsy at Rasoul Akram Hospital, Tehran, Iran, between 2023 and 2024. Ethical approval was obtained from the Institutional Review Board of Iran University of Medical Sciences (IR.IUMS.FMD.REC.1403.377), and all procedures adhered to the Declaration of Helsinki.

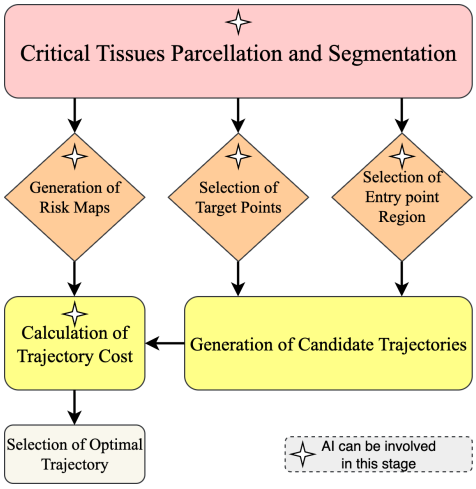

Fig. 1. Overall pipeline of automated path planning algorithms.

Imaging modality selection is critical for safe trajectory planning. Although time-of-flight MR angiography (TOF-MRA) facilitates arterial segmentation, it provides limited soft-tissue contrast and inadequately visualizes small, slow-flow, and venous vessels; particularly near tumor margins—potentially compromising safety in complex anatomical regions. In contrast, T1CE MRI offers superior soft-tissue contrast and enables concurrent visualization of arterial and venous vasculature, including small-caliber vessels essential for defining safe surgical corridors. However, non-specific enhancement of tumors and surrounding tissues introduces signal heterogeneity, making vessel segmentation challenging and rendering manual annotation labor-intensive and expertise-dependent.

Each MRI volume comprised 160 axial slices with 256 × 256 in-plane resolution and near-isotropic voxel spacing (≈ 1 mm; mean in-plane 0.969 mm, through-plane 1.000 mm). All scans were manually annotated using 3D Slicer by a board-certified functional neurosurgeon with more than 20 years of stereotactic experience, delineating cerebral vasculature (arterial and venous), tumor core boundaries, and cerebrospinal fluid spaces including the lateral, third, and fourth ventricles.

B. Image Preprocessing Pipeline

A standardized preprocessing pipeline was applied to optimize T1CE MRI volumes for DL–based segmentation (Fig. 2(a)). Non-brain tissues were removed using HD-BET [11], selected for its robustness to pathological anatomy, reducing computational load and false-positive detections from extracranial structures. Intensity inhomogeneities were corrected using N4 bias correction, a critical step for T1CE imaging, where spatially varying enhancement can obscure tissue boundaries.

Following skull stripping and bias correction, voxel intensities were linearly normalized to [0, 255] per volume, and

all scans were resampled to a fixed size of 192 × 192 × 128 using trilinear interpolation. Center-cropping or zero-padding was applied as needed to preserve anatomical proportions and focus learning on the central brain region containing the tumor and critical vasculature.

To improve generalization, data augmentation was applied during training only, including random vertical flipping along the y-axis (30% probability) and additive Gaussian noise (mean 0, SD 0.01) applied exclusively to background voxels (intensity ¡ 10). No augmentation was used for validation or testing.

C. DL-Based Segmentation of Critical Structures

Three critical anatomical structures—cerebral vasculature, tumor core, and CSF spaces, were segmented using a fusion of DL and semi-automatic methods (Fig. 2(b)).

Cerebral arteries and veins were segmented using NeuroVASC U-Net, a 3D patch-based CNN optimized for T1CE MRI [9]. The model follows a U-Net encoder–decoder architecture with residual connections and channel-wise attention, and was trained on $64\times64\times64$ patches using a combined Dice and binary cross-entropy loss. Training employed the Adam optimizer (learning rate 1×10^{-4}, batch size 4) with early stopping. On a held-out validation cohort (n = 28), the model achieved a dice similarity coefficient (DSC) score of 86.09% and precision of 88.41%, demonstrating reliable segmentation of both large vessels and fine perforators.

Tumor core segmentation was performed using a 3D Attention U-Net operating on full-resolution volumes $(192 \times 192 \times 128)$. Multi-scale attention improved boundary delineation in heterogeneously enhancing tumors, yielding a validation DSC score of 75.79%. The biopsy target point (P_{target}) was defined as the centroid of the largest connected tumor component after DBSCAN-based removal of spurious clusters (minimum size: 100 voxels).

CSF spaces were segmented semi-automatically based on their hypointense appearance on T1CE MRI. Initial masks were generated using intensity thresholding (threshold = 15) and morphological closing (spherical kernel, radius = 2 voxels), followed by expert manual review and correction to include missed regions such as the aqueduct of Sylvius and foramina of Monro.

D. Risk Map Construction and Obstacle Integration

To enable quantitative trajectory evaluation, segmented anatomical structures were integrated into a unified obstacle representation. Cerebral vasculature and CSF spaces were merged into a binary obstacle map, as traversal of either constitutes an absolute surgical contraindication (hemorrhage or CSF leakage/infection). Tumor masks were excluded, as the biopsy target is defined within the tumor core.

1) Euclidean Distance Transform: A 3D Euclidean distance transform (EDT) was computed on the obstacle map to generate a continuous distance field, where each voxel value $D(v)$ represents the shortest Euclidean distance to the nearest obstacle. The EDT was computed using the fast marching method with $\mathcal{O}(N)$ complexity and forms the basis for trajectory cost evaluation by penalizing paths that approach critical structures.

2) Probe Radius Consideration: To account for the finite biopsy needle diameter, a probe-aware distance metric was incorporated. Assuming a 6-mm probe ($\approx$ 6 voxels at 1-mm isotropic resolution), a spherical neighborhood of radius $r = 3$ voxels was evaluated at each trajectory voxel. The effective clearance distance $D_n(v)$ was defined as Eq. 1:

$$D_n(v) = \min_{u \in B_r(v)} D(u) \tag{1}$$

where $B_r(v)$ denotes the set of voxels within a sphere of radius r centered at v, and $D(u)$ is the EDT value at voxel u. Ensuring that the full probe cross-section maintains adequate clearance from obstacles. This formulation yields a risk map that more accurately reflects true surgical constraints than centerline-only evaluation.

E. Automated Skull Entry Point Identification

Automated identification of anatomically valid skull entry points is a key component of the proposed framework. Unlike conventional stereotactic workflows that rely on manually defined cranial ROIs (10–30 cm^2), which are time–consuming and operator–dependent, we adopt a geometry–based strategy that extracts the entire cranial surface as a dense set of candidate entry points and automatically restricts it to cosmetically and surgically acceptable regions above the hairline.

Candidate points are obtained by extracting a single–voxel–thick skull surface shell from skull–stripped T1CE MRI using classical morphological operations. Briefly, a binary brain mask is duplicated; one copy is isotropically dilated by one voxel (26–connectivity), both masks are cavity–filled, and their difference yields a thin shell representing the brain–skull interface. This process preserves full cranial geometry with uniform sampling, requires no manual annotation or learning–based inference, and executes in under 0.5 s per volume.

Anatomically invalid regions (facial skeleton, orbital roofs, nasal cavity, and posterior cervical areas) are removed using a patient–specific geometric filter defined by a separating hyperplane. Two landmarks are placed per patient in < 10 s: an anterior–superior forehead point $\mathbf{c}_1$ and a posterior–inferior upper–neck point $\mathbf{c}_2$, defining a sagittal–plane vector $\mathbf{v}_{yz} = (\Delta z, \Delta y)$. The normalized hyperplane normal $\mathbf{n}$ (Eq. 2) and the plane Π through the midpoint $\mathbf{m} = \frac{1}{2}(\mathbf{c}_1 + \mathbf{c}_2)$ (Eq. 3) define a valid entry domain as the superior half–space H^+ (Eq. 4).

$$\mathbf{n} = \frac{1}{\|\mathbf{v}_{yz}\|_2} \begin{bmatrix} 0 \\ -\Delta z \\ \Delta y \end{bmatrix} \tag{2}$$

$$\Pi = \left\{ \mathbf{x} \in \mathbb{R}^3 \mid \mathbf{n}^\top (\mathbf{x} - \mathbf{m}) = 0 \right\}. \tag{3}$$

$$H^+ = \left\{ \mathbf{x} \mid \mathbf{n}^\top (\mathbf{x} - \mathbf{m}) \geq 0 \right\}. \tag{4}$$

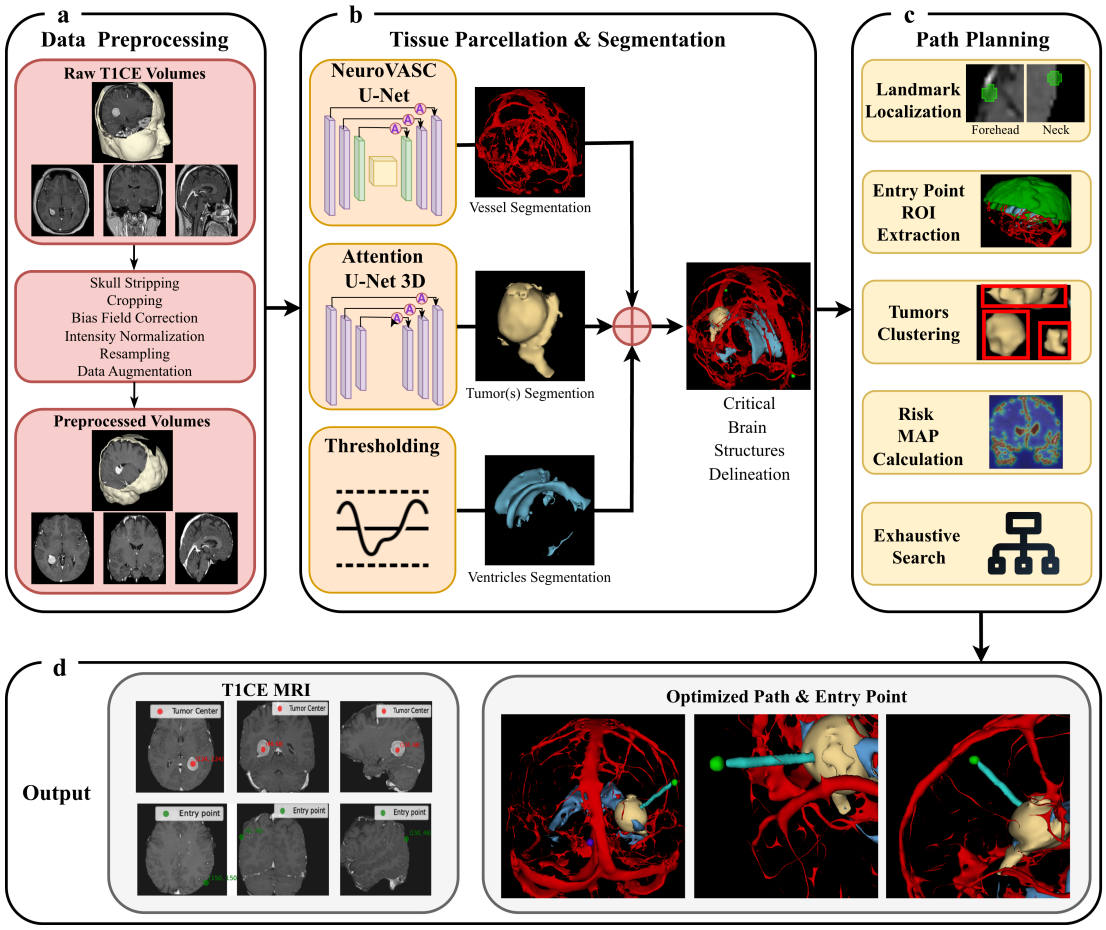

Fig. 2. Caption

This filtering yields an average of 27,051 valid entry voxels per patient (range: 18,342–35,786), substantially exceeding typical manual ROIs (500–2,000 voxels) while reducing manual effort. Because the hyperplane is derived from patient–specific landmarks, the method adapts to head–position variability during MRI acquisition and ensures consistent anatomical validity across subjects.

F. Deterministic Trajectory Optimization

To guarantee identification of the safest feasible biopsy trajectory with global optimality and full reproducibility, a deterministic exhaustive search strategy was employed. Unlike stochastic optimizers and sampling–based planners (e.g., genetic algorithms, simulated annealing, RRT*, PRM), which may converge to local minima or exhibit run–to–run variability, exhaustive enumeration evaluates all candidate trajectories in the discretized search space, making it well suited for preoperative planning.

For each valid skull entry point $P_{\text{entry}} \in H^+$ and tumor target point P_{target}, a straight–line trajectory was discretized using the three–dimensional Bresenham algorithm, yielding an ordered set of integer voxel coordinates L. This enables direct evaluation on the voxelized anatomical distance field. To mitigate discretization artifacts caused by integer rounding—particularly when the floating–point tumor centroid does not align with voxel centers—an error–correction step was applied near the target to ensure that the terminal voxel lies within the tumor core mask.

Trajectory evaluation was formulated as a probe–aware cost minimization problem that balances clearance from critical anatomy against intracerebral path length. The total trajectory cost is defined in Eq. 5:

$$J = \sum_{v \in L} \frac{1}{\left(D_n(v) + \varepsilon\right)^2} + \lambda \left\| P_{\text{entry}} - P_{\text{target}} \right\|_2, \qquad (5)$$

where $D_n(v)$ is the minimum distance to obstacles within a spherical neighborhood of radius $r = 3$ voxels, computed as defined in Eq. 6:

$$D_n(v) = \min_{u \in B_r(v)} D(u). \qquad (6)$$

Here, $D(u)$ denotes the Euclidean distance transform value at voxel u, and $\varepsilon = 10^{-6}$ ensures numerical stability near obstacle boundaries. The inverse–square term strongly penalizes proximity to critical structures, while the path–length term discourages excessively long trajectories, with $\lambda = 0.05$ empirically chosen to balance safety and invasiveness.

Hard geometric constraints were imposed to ensure surgical feasibility. Trajectories were discarded if their total length exceeded 90 mm or if the drilling angle relative to the local skull surface normal exceeded 15°. The drilling angle θ is computed as defined in Eq. 7:

$$\theta = \arccos\left(\frac{\mathbf{n}_{\text{skull}} \cdot (P_{\text{target}} - P_{\text{entry}})}{\|P_{\text{target}} - P_{\text{entry}}\|_2}\right), \tag{7}$$

where $\mathbf{n}_{\text{skull}}$ is the local skull surface normal at P_{entry}, estimated from the gradient of the skull shell mask.

To reduce computational overhead, collision detection was performed prior to cost evaluation. A trajectory was considered collision–free only if $D_n(v) > 0$ for all $v \in L$, ensuring no intersection with vascular or ventricular structures. The globally optimal trajectory is selected by exhaustive enumeration as defined in Eq. 8:

$$\tau^* = \arg\min_{P_{\text{entry}} \in H^+} J(P_{\text{entry}}, P_{\text{target}}), \tag{8}$$

subject to collision–free, length, and angle constraints. On average, approximately 27,000 candidate entry points were evaluated per patient, yielding a mean of 545 collision–free trajectories per tumor (range: 12–1,847). The median computation time was 10.54 ± 1.15 s per patient on consumer–grade hardware (NVIDIA RTX 4060, 8 GB VRAM).

1) Generation of Spatially Distinct Alternative Trajectories: To support intraoperative flexibility, the framework reports a ranked set of spatially distinct candidate trajectories in addition to the global optimum. After identifying the minimum–cost collision–free trajectory τ^*, additional candidates are selected from the remaining feasible set by enforcing a minimum Euclidean separation between skull entry points, as defined in Eq. 9:

$$\|P_{\text{entry}}^{(k)} - P_{\text{entry}}^{(j)}\|_2 \geq 10 \text{ mm}, \quad \forall j < k. \tag{9}$$

This constraint suppresses near–duplicate solutions and promotes genuinely distinct surgical corridors (e.g., anterior, lateral, posterior). Trajectories satisfying Eq. 9 are ranked using the same cost function, and the five lowest–cost solutions are retained. For the representative patient in Table II, trajectory costs ranged from 5.16 (global optimum) to 7.81 (fifth–ranked alternative).

2) Sequential Planning for Multiple Tumors: For patients with multiple intracranial lesions, trajectories are planned sequentially with dynamic updates of the obstacle map to prevent inter–trajectory interference. After selecting the optimal trajectory τ_1 for the first tumor, its centerline voxels—expanded by one probe radius—are added to the obstacle map, and the Euclidean distance transform is recomputed. The optimization process is then repeated to obtain τ_2, and iteratively applied for additional lesions, yielding a set of mutually non–intersecting trajectories $\{\tau_i\}$ that remain individually optimal under the updated constraints.

This sequential strategy avoids the computational burden of joint multi–lesion optimization while guaranteeing collision–free access. If no feasible trajectory exists after incorporating previously planned paths, the system flags tumor–specific infeasibility and suggests alternatives such as reducing probe diameter (2–3 mm) or reordering the biopsy sequence.

III. Results

A. Computational Performance and Planning Success

The automated trajectory–planning framework was implemented in Python and evaluated on consumer–grade hardware (NVIDIA RTX 4060, 8 GB VRAM). Validation on a cohort of 137 patients demonstrated fast and robust performance, with a mean planning time of 10.54 ± 1.15 s per patient. This represents an approximately tenfold reduction compared with reported manual planning times following segmentation (12–30 min).

Using a standard 6–mm diameter biopsy probe, the framework successfully identified collision–free trajectories in 135 of 137 patients, corresponding to a success rate of 98.5%. The two failures involved deep–seated tumors completely surrounded by dense vascular networks. In these cases, the system explicitly flagged trajectory infeasibility and recommended smaller probe diameters (2–3 mm), demonstrating clinically meaningful failure detection rather than silent optimization errors. For the trajectories selected as globally optimal, the mean entry-to-target Euclidean distance across the cohort was 44.4 ± 19.9 mm (median 41.1 mm; IQR 28.3–62.7 mm), corresponding to an approximate mean intracerebral path length of 36–38 mm after subtraction of typical scalp and skull traversal (≈6–8 mm). This is comparable to the 38.5 mm intracerebral length reported for computer-assisted biopsy planning by Marcus *et al.* [12] and is well within the imposed 90 mm hard constraint. Combined with the mean computation time of 10.54 ± 1.15 s per patient, these results indicate that the framework produces trajectories of clinically conventional length at a substantially reduced planning-time cost relative to manual workflows.

Automatic skull entry–zone generation produced an average of $27{,}051 \pm 3{,}842$ candidate entry points per patient, substantially exceeding the size of typical manually defined regions of interest (500–2,000 points). Despite this enlarged search space, computational efficiency was maintained through GPU–accelerated Euclidean distance transform computation.

B. Trajectory Safety, Optimality, and Multi–Path Planning

Across the entire cohort, the global average cost of the model–selected optimal trajectory was 5.91 ± 6.22, indicating consistent identification of low–risk corridors across diverse anatomical configurations. In a representative case (Fig. 2(c)), the system evaluated 27,573 candidate entry points, identified 545 valid non–colliding trajectories, and reported the five lowest–cost solutions with costs ranging from 5.16 to 7.81. A minimum spatial separation of 10 mm between skull entry points was enforced, yielding multiple clinically distinct and viable alternatives for intraoperative decision–making.

For cases involving multiple intracranial lesions, sequential planning with dynamic obstacle updating successfully generated non–intersecting trajectories when geometrically feasible.

Representative planning outputs were qualitatively reviewed by board–certified neurosurgeons at two institutions, who confirmed anatomical plausibility and adherence to standard stereotactic constraints, including entry angle $\leq 15^{\circ}$, path length ≤ 90 mm, and midline avoidance. Representative multi–tumor planning outcomes, including the number of valid collision–free paths and associated costs, are summarized in Table **??**. The table highlights both successful multi–trajectory planning and isolated cases in which no feasible trajectory could be identified under the specified probe diameter.

C. Probe-Diameter Sensitivity and Failure-Mode Analysis

To support direct comparison against current clinical practice and to characterise the dependence of trajectory feasibility on the probe-aware clearance envelope, the framework was re-executed across the entire cohort (n=137 patients, 402 tumor clusters) for three probe diameters (4, 5, and 6 mm). The exclusion radius modelled here corresponds to the physical biopsy cannula (typical outer diameter 1.8–2.5 mm in modern stereotactic practice [13], [14]) plus a configurable safety margin chosen by the surgeon to account for trajectory placement error and a vascular standoff distance comparable to the 3 mm margin commonly applied in stereo-EEG practice [15].

As summarised in Table II, at clinically realistic probe envelopes of 4 and 5 mm the framework identified at least one collision-free trajectory in 136 of 137 patients (patient-level success 99.3%) and for 396 of 402 tumor clusters (cluster-level success 98.5%). Increasing the probe envelope to 6 mm reduced patient-level feasibility to 92.7% (127/137) and cluster-level feasibility to 70.6% (284/402), reflecting the substantially stricter geometric clearance required around vasculature for a larger envelope. These results indicate that trajectory feasibility is highly sensitive to the configured clearance envelope, and that probe geometry, together with the surgeon-selected vascular safety margin, is a critical determinant of planning success that should be reported explicitly in any automated stereotactic biopsy workflow.

The 9 cases without a feasible trajectory at 6 mm and the 2 cases at 4–5 mm uniformly involved deep-seated lesions surrounded by dense vascular networks. In these instances, the framework explicitly flagged infeasibility rather than returning a low-quality plan, and reported the clearance-limiting structures, allowing the surgeon to consider a smaller-diameter needle (1.8–2.0 mm, e.g., [13]), a reduced vascular margin under supervisory review, or an alternative biopsy strategy. This explicit failure-detection behaviour is, to our knowledge, not reported by prior deterministic planners and is a deliberate design choice to maintain clinical transparency.

D. Expert Evaluation of Planned Trajectories

To further assess clinical relevance, the model–suggested optimal trajectory for each collision–free case was independently reviewed by a board–certified neurosurgeon and categorized as *optimal*, *near optimal*, or *not optimal* based on anatomical plausibility and surgical feasibility. Among the 135 collision–free cases, 101 trajectories (74.8%) were rated as optimal, 23 (17.0%) as near optimal, and the remaining 11 cases (8.1%) were classified as not optimal. This expert assessment demonstrates strong concordance between the automated optimization framework and neurosurgical judgment, supporting the clinical validity of the proposed approach.

IV. Discussion

This study presents a fully automated trajectory–planning framework integrating three advances: (1) DL–based segmentation of cerebral vasculature (NeuroVASC U–Net) and tumor core (3D Attention U–Net) directly from T1CE MRI, eliminating manual segmentation while capturing arterial and venous anatomy absent in TOF–MRA; (2) a geometric constraint that automatically defines anatomically valid skull entry zones; and (3) a deterministic exhaustive search that guarantees global optimality, supports five spatially distinct alternative trajectories, and enables sequential multi–tumor planning via dynamic obstacle updates.

A. Comparison with Existing Literature

Most prior automated planners rely on TOF–MRA for vessel visualization because of simplified arterial segmentation, despite its poor soft–tissue contrast, limited depiction of slow–flow vessels, and near–absence of venous anatomy; limitations that are particularly problematic in peritumoral regions. Studies using T1CE MRI required semi–automatic segmentation with expert correction [16], [17], typically adding 30–60 min per case. To our knowledge, this is the first study to apply validated DL–based vessel segmentation on T1CE MRI for stereotactic planning, enabling comprehensive vascular modeling without manual annotation.

Existing approaches depend on manual ROIs [17]–[19], generic skull templates [16], or exhaustive searches over the entire skull surface [12], [20]. Manual ROIs introduce operator dependence and restrict the search to 500–2,000 points, while full–surface searches require substantial computation (up to 3.3 min). The proposed geometric constraint removes manual ROI placement, increases the search space by approximately 13×, and retains computational efficiency ($<$ 11 s) through anatomically informed pruning.

Sampling–based planners (RRT, PRM) exhibit stochastic behavior and local optimality [21], while genetic algorithms reduce runtime [22], [23] but lack global optimality guarantees. Exhaustive search ensures optimality but is often computationally prohibitive without search–space reduction [12], [17], [24]. By constraining the entry zone, our framework achieves deterministic global optimality with competitive runtime (10.54 s) and reports multiple spatially diverse solutions, whereas most prior studies provide only a single trajectory.

To benchmark the proposed framework against current clinical practice more directly, we contextualise our reported computation time and planning success rate alongside published manual and computer-assisted comparators. Manual stereotactic trajectory planning has been reported to require on the order of 15 minutes per electrode in stereo-EEG planning, accumulating to 2–3 hours per multi-electrode case [16], while

TABLE I

REPRESENTATIVE MULTI-TUMOR PLANNING OUTCOMES (PROBE ENVELOPE: 5 MM). $|E|$ DENOTES THE NUMBER OF ANATOMICALLY VALID SKULL ENTRY VOXELS CONSIDERED. "VALID PATHS" DENOTES THE NUMBER OF COLLISION-FREE CANDIDATE TRAJECTORIES IDENTIFIED BY EXHAUSTIVE SEARCH FOR THAT TUMOR CLUSTER UNDER THE CURRENT OBSTACLE MAP. CASES WITH NO FEASIBLE TRAJECTORY ARE REPORTED EXPLICITLY TO DEMONSTRATE TRANSPARENT FAILURE DETECTION.

Patient	Tumor	Target (x, y, z)	$\lvert E\rvert$	Best entry (x, y, z)	Valid paths	Cost / Outcome
26	1	[115, 105, 52]	40866	[144, 91, 107]	962	4.6441 / Success
26	2	[128, 91, 86]	40866	[146, 68, 102]	1093	2.9192 / Success
5	1	[51, 122, 39]	22687	[71, 162, 98]	231	5.3575 / Success
5	2	[56, 53, 62]	22687	[42, 63, 92]	219	5.1504 / Success
5	3	[84, 96, 93]	22687	[73, 91, 122]	352	3.4312 / Success
5	4	[59, 56, 86]	22687	[50, 52, 96]	51	4.8987 / Success
5	5	[75, 165, 42]	22687	[84, 178, 57]	141	1.5427 / Success
67	1	[59, 62, 37]	28616	[92, 132, 118]	19	14.5604 / Success
67	2	[117, 99, 72]	28616	[150, 105, 103]	415	4.2234 / Success
67	3	[121, 83, 25]	28616	–	0	– / **(no valid path)**
76	1	[80, 75, 96]	22435	–	0	– / **(no valid path)**
76	2	[86, 74, 87]	22435	–	0	– / **(no valid path)**
76	3	[105, 75, 107]	22435	[121, 89, 124]	765	8.4382 / Success
79	1	[97, 76, 63]	21220	[154, 125, 83]	26	8.4382 / Success
79	2	[108, 74, 83]	21220	–	0	– / **(no valid path)**
90	1	[42, 65, 33]	31458	–	0	– / **(no valid path)**
90	2	[53, 74, 38]	31458	[87, 20, 86]	1	14.4576 / Success
90	3	[55, 107, 28]	31458	[28, 127, 47]	10	8.4395 / Success
90	4	[71, 101, 76]	31458	[56, 100, 105]	1234	2.4990 / Success
90	5	[119, 140, 31]	31458	[141, 170, 23]	216	5.8065 / Success
101	1	[47, 100, 76]	35213	[26, 11, 75]	1213	1.8627 / Success
101	2	[91, 131, 54]	35213	[68, 175, 87]	1093	4.4702 / Success
103	1	[49, 123, 80]	36348	[120, 174, 20]	296	2.8700 / Success
103	2	[63, 92, 80]	36348	[71, 112, 118]	1151	3.1525 / Success

TABLE II

TRAJECTORY FEASIBILITY BY PROBE DIAMETER ACROSS THE FULL COHORT (137 PATIENTS, 402 TUMOR CLUSTERS), REPORTED AT BOTH THE PATIENT LEVEL (AT LEAST ONE FEASIBLE TRAJECTORY PER PATIENT) AND THE CLUSTER LEVEL (A FEASIBLE TRAJECTORY FOR EACH INDIVIDUAL TUMOR CLUSTER).

Probe diameter	Patients with ≥ 1 feasible path	Patient-level success (%)	Tumor clusters (n)	Clusters successfully planned	Cluster-level success (%)
4 mm	136	99.3	402	396	98.5
5 mm	136	99.3	402	396	98.5
6 mm	127	92.7	402	284	70.6

manual stereotactic biopsy planning is typically performed in 12–30 minutes per case [8], [12]. In the comparative pilot study by Marcus et al. [12], computer-assisted planning produced trajectories with shorter intracerebral length (38.5 mm vs. 43.5 mm; p=0.01), reduced risk score (0.27 vs. 0.52; p=0.03), and a median planning time of 8 minutes. The present framework achieves a comparable risk-aware optimisation in 10.54 ± 1.15 s per patient, while additionally considering a substantially larger entry-point search space ($\approx$ 27,000 candidates vs. the 500–2,000 candidates typical of manually defined regions of interest [12], [17]) and returning five spatially distinct alternative trajectories rather than a single optimum. We emphasise, in line with the recent IDEAL-D Stage-0 categorisation by Starup-Hansen et al. [8], that these benchmarks describe surrogate planning metrics rather than intraoperative outcomes; the present comparison is therefore positioned as a workflow- and safety-margin benchmark, and a head-to-head, blinded comparison against the manually executed trajectories used in this retrospective cohort is identified as the necessary next step. These benchmarks are summarised in Table III.

A recent systematic review [8] classified all published automated planners as IDEAL–D Stage 0 [25], citing the absence of prospective trials, ground–truth comparison, and outcome validation. Our work also represents Stage 0 validation but advances the field through evaluation on a prospectively collected cohort (n=137), high planning success (98.5%), and qualitative expert acceptance. Nonetheless, prospective blinded outcome studies remain necessary for clinical translation.

B. Clinical Implications, Future Directions, and Limitations

This study has several limitations that also define key priorities for future work.

Retrospective design and need for prospective validation. The current evaluation is retrospective, and expert review, although supportive of clinical plausibility, cannot replace prospective workflow-integrated validation. In line with the IDEAL-D framework [25] and prior classification of AI-based biopsy planning as IDEAL-D Stage 0 [8], the next step should be a prospective IDEAL-D Stage 1 study. This should include blinded comparison of automated and manual trajectories, prospective assessment of diagnostic yield and hemorrhagic complications, and multi-centre external validation. Because this cohort consisted of patients who had already undergone biopsy, important workflow metrics such as intraoperative timing, recorded entry-point coordinates, and manual planning times were not consistently available; these should be collected prospectively to allow direct quantitative comparison with manual planning.

TABLE III

COMPARISON OF PLANNING TIME, TRAJECTORY LENGTH, AND MULTI-TRAJECTORY OUTPUT OF THE PROPOSED FRAMEWORK AGAINST REPRESENTATIVE PUBLISHED MANUAL AND COMPUTER-ASSISTED STEREOTACTIC PLANNING METHODS. NR = NOT REPORTED. SEEG = STEREO-ELECTROENCEPHALOGRAPHY. BIOPSY TIMES FOR THE SPARKS AND VAKHARIA SEEG PLANNERS ARE REPORTED PER-ELECTRODE AND PER-CASE RESPECTIVELY; THE PROPOSED FRAMEWORK IS BENCHMARKED AGAINST THE MOST DIRECTLY COMPARABLE BIOPSY WORK [12].

Study (year)	Procedure	Method	Planning time	Trajectory length (mm)	Multi-traj. output
Trope *et al.* [17] (2015)	Biopsy	Automated, manual ROI	NR	NR	Single
Sparks *et al.* [16] (2017)	SEEG	MTP / EpiNav (auto)	<1 min/case	NR	Yes
Vakharia *et al.* [20] (2018)	SEEG	EpiNav, external val.	rapid	NR	Yes
Marcus *et al.* [12] (2020)	Biopsy	Manual planning	12–30 min	43.5	Single
Marcus *et al.* [12] (2020)	Biopsy	SurgiNav (CAP)	8 min (median)	38.5	Ranked
Proposed	**Biopsy**	**Deterministic exhaustive**	10.54 ± 1.15 **s**	$\approx$ 36–38 **(intracerebral)**	**Yes (5 distinct)**

Rigid-anatomy assumption and intraoperative brain shift. The framework currently assumes that preoperative MRI adequately represents intraoperative anatomy. Although brain shift in open craniotomy may reach 1–2 cm at the cortical surface and exceed 3 mm at deep tumor margins in most cases [26], the magnitude is generally much smaller in stereotactic biopsy. In a study of 44 burr-hole deep-brain stimulation procedures using interventional MRI, Ivan *et al.* [27] reported a mean displacement of 0.6 mm at the ipsilateral deep target, with shifts greater than 2 mm in only 9% of patients. In the present framework, this residual uncertainty can be partly mitigated by increasing the configurable probe-radius parameter, allowing a practical safety margin along the trajectory. While this does not substitute for intraoperative deformable registration, it offers a clinically reasonable mitigation for biopsy planning. Future work should explore integration of deformable registration and biomechanical modelling for real-time trajectory adjustment, as previously investigated in deep-brain stimulation [28].

Probe geometry and smaller-needle alternatives. Trajectory feasibility was highly sensitive to the selected probe envelope, as noted in Section III-C. The complete failures observed at 4–5 mm, along with additional infeasible cases at 6 mm, all involved deeply vascularised lesions. Smaller-diameter cannulas (e.g., 1.8–2.0 mm), which can preserve diagnostic yield while reducing hemorrhage risk [13], [14], are already compatible with the framework through its configurable probe parameter. Future work should also investigate steerable needles and concentric-tube robots to expand access to difficult targets [29], while recognising the added mechanical and regulatory complexity of these approaches.

Segmentation uncertainty. The current pipeline treats deep-learning-derived vessel, tumor, and ventricle masks as deterministic obstacles. However, recent methods for uncertainty quantification in medical image segmentation, including Monte Carlo dropout [30] and test-time augmentation [31], could support uncertainty-aware planning. For example, voxel-wise uncertainty could be incorporated into the cost function so that trajectories passing through uncertain regions are penalised and adaptive safety margins are applied where anatomy is ambiguous. This should be a priority for future development.

Additional limitations include the single-centre cohort, which limits generalisability and supports the need for multi-centre validation; the restriction to straight trajectories, which reduces reachability in deeply vascularised regions; and the use of fixed cost-function weights, which do not yet reflect surgeon-specific preferences.

C. Conclusion

This study presents a fully automated trajectory planning framework achieving 98.5% success rate and 10.54-second computation time through integration of DL-based T1CE vessel segmentation (DSC=86.09%), geometric hairline-based entry zone definition (27,051 candidate points), and deterministic exhaustive search with multi-path reporting. Comparison with literature demonstrates advantages over prior approaches: comprehensive vascular visualization (vs. TOF-MRA limitations), automated entry zone definition (vs. manual ROI placement), deterministic global optimality (vs. stochastic methods), and multi-path output (vs. single-trajectory systems). The framework addresses critical workflow bottlenecks (10× faster than manual planning) while maintaining clinical feasibility, positioning it for prospective validation (IDEAL-D Stage 1) with blinded outcome assessment—a necessary step toward regulatory approval and clinical translation. Future work will integrate intraoperative imaging for brain shift compensation, extend to curved trajectory planning for anatomically complex cases, and implement uncertainty-aware cost functions to enhance safety guarantees.

REFERENCES

[1] D. Thomas, "Stereotactic biopsies for tumors: indications, limits, diagnosis with histopathology and other laboratory techniques," in *Practical Handbook of Neurosurgery: From Leading Neurosurgeons*. Springer, 2009, pp. 529–547.

[2] P. Kickingereder, P. Willeit, T. Simon, and M. I. Ruge, "Diagnostic value and safety of stereotactic biopsy for brainstem tumors: a systematic review and meta-analysis of 1480 cases," *Neurosurgery*, vol. 72, no. 6, pp. 873–882, 2013.

[3] F. Jozsa, C. Gaier, Y. Ma, N. Kitchen, A. McEvoy, A. Miserocchi, G. Samandouras, H. Sethi, L. Thorne, C. Hill *et al.*, "Safety and efficacy of brain biopsy: Results from a single institution retrospective cohort study," *Brain and Spine*, vol. 3, p. 101763, 2023.

[4] M. Zanello, A. Roux, C. Debacker, S. Peeters, M. Edjlali-Goujon, F. Dhermain, E. Dezamis, C. Oppenheim, E. Lechapt-Zalcman, M. Harislur *et al.*, "Postoperative intracerebral haematomas following stereotactic biopsies: Poor planning or poor execution?" *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 17, no. 2, p. e2211, 2021.

[5] N. Sugii, M. Matsuda, T. Tsurubuchi, and E. Ishikawa, "Hemorrhagic complications after brain tumor biopsy: risk-reduction strategies based on safer biopsy targets and techniques," *World neurosurgery*, vol. 176, pp. e254–e264, 2023.

[6] S. W. Hong, P. D. Dao, K. W. Chang, H. H. Jung, and J. W. Chang, "Minimizing hemorrhage complications in deep brain stimulation surgery-the impact of imaging modalities and trajectory planning," *Journal of Korean Neurosurgical Society*, vol. 68, no. 5, pp. 600–608, 2025.

[7] R. Monfaredi, A. Concepcion-Gonzalez, J. Acosta Julbe, E. Fischer, G. Hernandez-Herrera, K. Cleary, and C. Oluigbo, "Automatic path-planning techniques for minimally invasive stereotactic neurosurgical procedures—a systematic review," *Sensors*, vol. 24, no. 16, p. 5238, 2024.

[8] J. Starup-Hansen, S. C. Williams, J. P. Funnell, J. G. Hanrahan, S. Islam, A. Al-Mohammad, and C. S. Hill, "Optimising trajectory planning for stereotactic brain tumour biopsy using artificial intelligence: a systematic review of the literature," *British Journal of Neurosurgery*, vol. 39, no. 2, pp. 163–172, 2025.

[9] M. J. Vayeghan, N. Delfan, M. T. Masouleh, M. P. Rizi, and B. Moshiri, "Neurovascu-net: A unified multi-scale and cross-domain adaptive feature fusion u-net for precise 3d segmentation of brain vessels in contrast-enhanced t1 mri," *arXiv preprint arXiv:2511.18422*, 2025.

[10] S. F. Abbas, N. T. Duc, Y. Song, K. Kim, E. Srivastava, and B. Lee, "Cv-attention unet: Attention-based unet for 3d cerebrovascular segmentation of enhanced tof-mra images," *arXiv preprint arXiv:2311.10224*, 2023.

[11] F. Isensee, M. Schell, I. Pflueger, G. Brugnara, D. Bonekamp, U. Neuberger, A. Wick, H.-P. Schlemmer, S. Heiland, W. Wick *et al.*, "Automated brain extraction of multisequence mri using artificial neural networks," *Human brain mapping*, vol. 40, no. 17, pp. 4952–4964, 2019.

[12] H. J. Marcus, V. N. Vakharia, R. Sparks, R. Rodionov, N. Kitchen, A. W. McEvoy, A. Miserocchi, L. Thorne, S. Ourselin, and J. S. Duncan, "Computer-assisted versus manual planning for stereotactic brain biopsy: a retrospective comparative pilot study," *Operative Neurosurgery*, vol. 18, no. 4, pp. 417–422, 2020.

[13] J. Yuen, C. X. Zhu, D. T. Chan, R. Y. Ng, W. Nia, W. S. Poon, H. K. Ng, V. C. Mok, L. K. Wong, T. C. Cheung *et al.*, "A sequential comparison on the risk of haemorrhage with different sizes of biopsy needles for stereotactic brain biopsy," *Stereotactic and functional neurosurgery*, vol. 92, no. 3, pp. 160–169, 2014.

[14] S. K. Akshulakov, T. T. Kerimbayev, M. Y. Biryuchkov, Y. A. Urunbayev, D. S. Farhadi, and V. A. Byvaltsev, "Current trends for improving safety of stereotactic brain biopsies: advanced optical methods for vessel avoidance and tumor detection," *Frontiers in oncology*, vol. 9, p. 947, 2019.

[15] F. Cardinale, M. Cossu, L. Castana, G. Casaceli, M. P. Schiariti, A. Miserocchi, D. Fuschillo, A. Moscato, C. Caborni, G. Arnulfo *et al.*, "Stereoelectroencephalography: surgical methodology, safety, and stereotactic application accuracy in 500 procedures," *Neurosurgery*, vol. 72, no. 3, pp. 353–366, 2013.

[16] R. Sparks, G. Zombori, R. Rodionov, M. Nowell, S. B. Vos, M. A. Zuluaga, B. Diehl, T. Wehner, A. Miserocchi, A. W. McEvoy *et al.*, "Automated multiple trajectory planning algorithm for the placement of stereo-electroencephalography (seeg) electrodes in epilepsy treatment," *International journal of computer assisted radiology and surgery*, vol. 12, no. 1, pp. 123–136, 2017.

[17] M. Trope, R. R. Shamir, L. Joskowicz, Z. Medress, G. Rosenthal, A. Mayer, N. Levin, A. Bick, and Y. Shoshan, "The role of automatic computer-aided surgical trajectory planning in improving the expected safety of stereotactic neurosurgery," *International journal of computer assisted radiology and surgery*, vol. 10, no. 7, pp. 1127–1140, 2015.

[18] R. R. Shamir, L. Joskowicz, I. Tamir, E. Dabool, L. Pertman, A. Ben-Ami, and Y. Shoshan, "Reduced risk trajectory planning in image-guided keyhole neurosurgery," *Medical physics*, vol. 39, no. 5, pp. 2885–2895, 2012.

[19] Y. Liu, P. E. Konrad, J. S. Neimat, S. B. Tatter, H. Yu, R. D. Datteri, B. A. Landman, J. H. Noble, S. Pallavaram, B. M. Dawant *et al.*, "Multisurgeon, multisite validation of a trajectory planning algorithm for deep brain stimulation procedures," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 9, pp. 2479–2487, 2014.

[20] V. N. Vakharia, R. Sparks, R. Rodionov, S. B. Vos, C. Dorfer, J. Miller, D. Nilsson, M. Tisdall, S. Wolfsberger, A. W. McEvoy *et al.*, "Computer-assisted planning for the insertion of stereoelectroencephalography electrodes for the investigation of drug-resistant focal epilepsy: an external validation study," *Journal of neurosurgery*, vol. 130, no. 2, pp. 601–610, 2018.

[21] R. Zelmann, S. Bériault, M. M. Marinho, K. Mok, J. A. Hall, N. Guizard, C. Haegelen, A. Olivier, G. B. Pike, and D. L. Collins, "Improving recorded volume in mesial temporal lobe by optimizing stereotactic intracranial electrode implantation planning," *International journal of computer assisted radiology and surgery*, vol. 10, no. 10, pp. 1599–1615, 2015.

[22] I. Villanueva-Naquid, C. Soubervielle-Montalvo, R. M. Aguilar-Ponce, S. Tovar-Arriaga, J. C. Cuevas-Tello, C. A. Puente-Montejano, M. Mejia-Carlos, and J. G. Torres-Corzo, "Risk assessment methodology for trajectory planning in keyhole neurosurgery using genetic algorithms," *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 16, no. 2, p. e2060, 2020.

[23] S. Peikert, C. Kunz, N. Fischer, M. Hlaváč, A. Pala, M. Schneider, and F. Mathis-Ullrich, "Automated linear and non-linear path planning for neurosurgical interventions," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 7731–7737.

[24] A. De León-Cuevas, S. Tovar-Arriaga, A. González-Gutiérrez, and M. A. Aceves-Fernández, "Risk map generation for keyhole neurosurgery using fuzzy logic for trajectory evaluation," *Neurocomputing*, vol. 233, pp. 81–89, 2017.

[25] A. Sedrakyan, B. Campbell, J. G. Merino, R. Kuntz, A. Hirst, and P. McCulloch, "Ideal-d: a rational framework for evaluating and regulating the use of medical devices," *Bmj*, vol. 353, 2016.

[26] C. Nimsky, O. Ganslandt, S. Cerny, P. Hastreiter, G. Greiner, and R. Fahlbusch, "Quantification of, visualization of, and compensation for brain shift using intraoperative magnetic resonance imaging," *Neurosurgery*, vol. 47, no. 5, pp. 1070–1080, 2000.

[27] M. E. Ivan, J. Yarlagadda, A. P. Saxena, A. J. Martin, P. A. Starr, W. K. Sootsman, and P. S. Larson, "Brain shift during bur hole–based procedures using interventional mri," *Journal of neurosurgery*, vol. 121, no. 1, pp. 149–160, 2014.

[28] N. Hamzé, A. Bilger, C. Duriez, S. Cotin, and C. Essert, "Anticipation of brain shift in deep brain stimulation automatic planning," in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2015, pp. 3635–3638.

[29] P. Sauerteig, M. K. Hoffmann, J. Mühlenhoff, G. Miccoli, D. Keiner, S. Urbschat, J. Oertel, T. Sattel, K. Flaßkamp, and K. Worthmann, "Optimal path planning for stereotactic neurosurgery based on an elastostatic cannula model," *IFAC-PapersOnLine*, vol. 55, no. 20, pp. 600–605, 2022.

[30] T. Nair, D. Precup, D. L. Arnold, and T. Arbel, "Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation," *Medical image analysis*, vol. 59, p. 101557, 2020.

[31] G. Wang, W. Li, M. Aertsen, J. Deprest, S. Ourselin, and T. Vercauteren, "Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks," *Neurocomputing*, vol. 338, pp. 34–45, 2019.