Towards Versatile and Safe Multi-Electrode Path Planning for SEEG Implantation

Ziqiao Qu[1,2†], Yangyang Shi[5†], Yuanqing Wang[3], Linxia Xiao[4(✉)], and Weixin Si[2(✉)]

[1]Southern University of Science and Technology, China
[2]Shenzhen University of Advanced Technology, China
[3]The Second People's Hospital of Shenzhen, China
[4]Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China
[5]School of Computer Science and Engineering, Central South University, China
siweixin@suat-sz.edu.cn, lx.xiao@siat.ac.cn

Abstract. Stereo-electroencephalography (SEEG) enables minimally invasive recording from deep and sulcal brain regions to localize epileptogenic networks for epilepsy surgery. However, traditional multiple SEEG electrodes planning remains largely manual, requiring experts to iteratively pick targets and search safe implantation paths while avoiding vessels, sulci, and other critical structures. To tackle this issue, we propose a coverage-aware planning method that integrates joint target localization and cooperative multi-electrode trajectory optimization into a unified reinforcement learning framework. Specifically, we introduce a capsule-based coverage model that represents functional recording sensitivity as a continuous region around each trajectory and maximizes overall region-of-interest (ROI) coverage through one-shot decision making. To ensure safety, we incorporate multi-layer constraints and inter-electrode spacing into deterministic projection with fixed constraint checking, enabling constraint-satisfying planning without manual penalty tuning. To validate the proposed method, we conducted planning on 8 retrospective temporal lobe epilepsy cases (84 electrodes) with CT/MRI-derived anatomical segmentations, achieving an 82% improvement in mean ROI coverage and an 88% increase in minimum inter-electrode spacing compared with existing methods while maintaining an 88% collision-free rate, supporting a versatile and safe planning framework that can adapt to diverse epilepsy treatment requirements and facilitate clinical deployment. Code is available at `https://github.com/Joe-Galaxy/safe-seeg-multielectrode-planning`.

Keywords: Multi-electrode planning · Trajectory optimization · Reinforcement learning · Coverage-aware planning

† Equal contribution.
(✉) Corresponding authors.

1 Introduction

Drug-resistant focal epilepsy affects a substantial fraction of patients with epilepsy and is associated with increased morbidity and mortality [1]. For selected candidates, epilepsy surgery can improve seizure control and quality of life, but requires precise localization of epileptogenic networks [2]. Stereo-electroencephalography (SEEG) supports this presurgical evaluation by implanting multiple depth electrodes to sample deep and sulcal structures [3], with complementary roles compared to subdural grids and strips [4]. Despite its clinical value, SEEG planning remains challenging because teams must jointly determine targets and safe paths that avoid critical anatomy [5] while maintaining operating-room accuracy [6, 7].

Recent computer-assisted planning (CAP) tools aim to reduce the manual workload and standardize trajectory design. Dasgupta *et al.* [8] enhance CAP with center-specific spatial priors to improve safety and reduce manual adjustments, while immersive visualization systems such as PreVISE provide virtual-reality support for interactive SEEG planning [9]. Beyond geometric safety, Dessert *et al.* [10] optimize patient-specific recording sensitivity as an objective for SEEG implantation. In parallel, personalized network modeling continues to improve hypothesis-driven targeting and interpretation [11]. However, existing approaches typically address target selection and trajectory optimization in separate stages or rely on soft penalties that may leave constraint violations when planning multiple electrodes simultaneously.

We formulate SEEG planning as sequential patient-level decision making with a one-shot action for each rigid electrode. Given patient-specific meshes and region-of-interest (ROI) candidates, the policy proposes a target refinement and insertion direction at each step, followed by deterministic projection and fixed checking under anatomical and inter-electrode constraints. Clinically, this action corresponds to the two decisions required for one rigid SEEG electrode: where the contact-bearing segment should sample the ROI and which insertion direction should connect it to a skull entry point. Accepted trajectories then become spacing and collision constraints for later electrodes, making the problem a coupled patient-level plan rather than a collection of independent trajectory regressions. This coupling is important because a target-direction pair that is acceptable in isolation may become undesirable after other electrodes have occupied nearby corridors.

The main contributions of this work are as follows.

- We formulate multi-electrode SEEG planning as reinforcement learning over patient-specific 3D meshes, jointly optimizing target locations and insertion trajectories.
- We introduce capsule-based coverage to refine candidate targets toward high-sensitivity regions while selecting straight insertion directions through per-electrode one-shot actions.
- We combine deterministic projection with fixed anatomical/inter-electrode constraint checking, and validate against CAP and sensitivity-based methods on retrospective temporal lobe epilepsy cases.

2 Related Work

Target localization and coverage modeling. Target selection in SEEG is driven by clinical hypotheses about the epileptogenic network, and recent quantitative pipelines aim to strengthen this process with data-driven evidence. For example, Wang *et al.* [11] propose a personalized modeling workflow to delineate epileptogenic networks from multimodal data. From the planning perspective, Dessert *et al.* [10] explicitly optimize patient-specific recording sensitivity to prioritize trajectories that better sample the intended regions. Multi-center studies further highlight that targeting conventions and implantation strategies vary substantially across institutions [12]. In practice, multi-electrode SEEG planning revolves around two key aspects: target localization and multi-electrode trajectory planning.

Multi-electrode trajectory planning under safety constraints. CAP systems compute collision-free trajectories by incorporating anatomical constraints and clinician-defined preferences, and can be enhanced with priors learned from historical implantations [8]. Spiegler *et al.* [9] introduce PreVISE to support interactive, immersive multi-trajectory planning in virtual reality. Optimization-based methods have also been explored for multi-electrode planning, including case-based formulations [13] and temporal lobe implantation strategies [14]. Beyond SEEG, multi-constrained trajectory planning has been studied in related interventions such as multi-needle thermal ablation [15], and safe reinforcement learning has been used to enforce hard constraints in screw trajectory planning [16]. Nevertheless, existing SEEG planners often decouple target selection from trajectory optimization and may not provide a unified mechanism to enforce inter-electrode constraints while optimizing coverage-driven objectives for all electrodes.

3 Method

3.1 Data processing and anatomical modeling

We retrospectively collected CT (post-implant) and MRI (pre-implant) volumes from 8 temporal lobe epilepsy cases. Each case was segmented into (i) gray matter and cortex surface, (ii) skull, (iii) vasculature, and (iv) deep structures (e.g., ventricles and brainstem). Vessels were segmented from available vascular imaging (e.g., MR angiography/venography or CT angiography) and converted into 3D meshes [5]. All structures were represented as triangle meshes, downsampled to 10k points, and converted to signed distance fields (SDFs) for vessel, deep-structure, cortex, and skull distance queries. For each electrode k, Ω_k denotes the segmented ROI mesh points, and a cone corridor from the ROI center constrains candidate entry locations. The SDF representation converts patient-specific vessel, deep-structure, skull, and cortex boundaries into fast geometric margin queries, while the cone corridor encodes clinically preferred burr-hole directions without replacing surgeon-defined hypotheses. Non-entry intersections

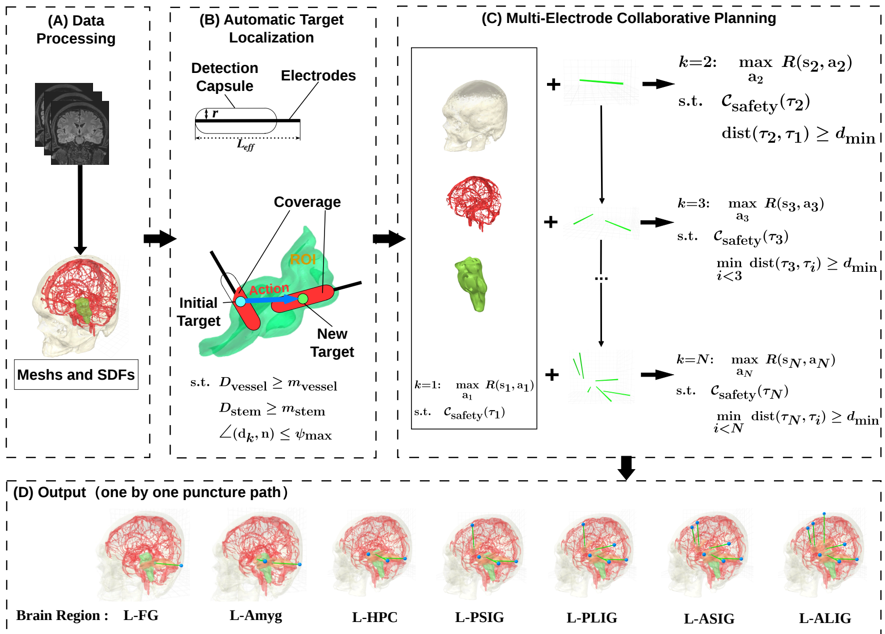

Fig. 1. Overview. (A) Patient-specific meshes are converted to signed distance fields (SDFs). (B) Capsule coverage refines ROI targets. (C) A one-shot RL policy proposes target offsets and directions; deterministic projection obtains candidate skull entries, and fixed checking validates anatomical and inter-electrode constraints. (D) The output is a checked multi-electrode trajectory set. R denotes the checked reward.

with cortex/skull are treated as forbidden events, whereas the selected skull entry is evaluated with the local skull normal for entry-angle checking.

3.2 MDP-guided one-shot multi-electrode implantation planning

We plan a set of electrodes indexed by $k = 1, \ldots, K$ (Fig. 1), where each electrode must (i) cover its ROI and (ii) satisfy strict anatomical safety constraints. We formulate patient-level planning as $M = (S, A, T, R, \gamma)$ with horizon K and $\gamma = 1$. At step k, the state contains anatomy, ROI/corridor, target-reference geometry, SDF distances, and accepted trajectories; the 5D action proposes one rigid electrode; the transition constructs/checks its straight trajectory and appends accepted paths as later spacing/collision constraints. Thus, "one-shot" means one electrode proposal, while the patient plan is produced by the sequential K-step process. This formulation follows the clinical workflow more closely than treating each implanted electrode as an independent sample. The policy observes previously accepted trajectories so that later proposals can account for inter-electrode spacing and collision rules, and the checked reward couples ROI

coverage with path preferences after geometric validation. PPO is used because no ground-truth target-direction labels are available for these coupled decisions.
Action and trajectory construction. As illustrated in Fig. 1(B), the policy outputs a 5D action $\mathbf{a}_k = (\Delta\mathbf{t}_k, \theta_k, \phi_k)$, where $\Delta\mathbf{t}_k \in \mathbb{R}^3$ refines the candidate target and (θ_k, ϕ_k) parameterize the insertion direction. We set $\mathbf{t}_k = \mathbf{t}_k^{(0)} + \Delta\mathbf{t}_k$ and convert angles to a unit direction $\mathbf{d}_k$. Ray-casting from $\mathbf{t}_k$ along $-\mathbf{d}_k$ gives the skull entry $\mathbf{e}_k$ and straight trajectory τ_k, matching the clinical output of target, burr-hole entry, and rigid shaft with workflow preferences encoded geometrically.
Capsule-based target localization via coverage. We model the recording sensitivity of a straight trajectory as a capsule of radius r (Fig. 1(B)), abstracting the contact-bearing shaft's effective sampling range; r follows electrode specifications or clinician-defined range and does not replace patient-specific biophysical models. As illustrated in Fig. 1(B), we define the capsule coverage of ROI mesh Ω_k as

$$c_k(\tau_k) = \frac{1}{|\Omega_k|} \sum_{\mathbf{v} \in \Omega_k} 1[d(\mathbf{v}, \tau_k) \le r], \tag{1}$$

where $d(\mathbf{v}, \tau_k)$ is the Euclidean distance from point $\mathbf{v}$ to the line segment of τ_k. Maximizing c_k encourages the target refinement $\Delta\mathbf{t}_k$ to move toward regions that are better covered by the trajectory, thereby achieving coverage-aware target localization.
Deterministic projection with multi-constraint checking. The RL policy selects ROI subregions, searches the 5D target-direction space, and learns geometry preferences from distance features, accepted trajectories, and checked rewards. Projection $\Pi(\cdot)$ only extends the proposed shaft to the skull entry and performs fixed accept/reject checking before reward evaluation; it is not a shortest-path planner and does not optimize capsule coverage. Given the policy proposal, checking validates vessel/deep-structure distance, forbidden non-entry cortex/skull intersections, path length, entry angle, corridor membership, and spacing to previously accepted electrodes. Only accepted trajectories are appended to the patient state and used as constraints for subsequent electrodes. Vessel, non-entry cortex/skull, and deep-structure distances use precomputed SDFs [18]. For an accepted electrode k, the checked trajectory $\hat{\tau}_k = \Pi(\tau_k)$ must satisfy:

$$\begin{aligned} &d(\hat{\tau}_k, \mathcal{V}) \ge \varepsilon_{\text{vessel}}, \quad d(\hat{\tau}_k, \mathcal{C}_{\text{non}}) \ge \varepsilon_{\text{cortex}}, \quad d(\hat{\tau}_k, \mathcal{B}) \ge \varepsilon_{\text{deep}}, \\ &L(\hat{\tau}_k) \le L_{\max}, \quad \angle(\hat{\tau}_k, \mathbf{n}_{\text{skull}}) \le \alpha_{\max}, \quad d(\hat{\tau}_k, \hat{\tau}_j) \ge \varepsilon_{\text{IE}}, \ \forall j < k, \end{aligned} \tag{2}$$

Here $\mathcal{V}$, $\mathcal{C}_{\text{non}}$, and $\mathcal{B}$ denote vessels, forbidden non-entry cortex/skull intersections, and deep structures; $\mathbf{n}_{\text{skull}}$ is the outward skull normal at entry, and ε_{IE} enforces inter-electrode spacing.
Training objective. Because Eq. (2) is checked before reward evaluation, the reward focuses on coverage and mild secondary preferences among accepted trajectories: $R_k = c_k(\hat{\tau}_k) - \lambda_L \frac{L(\hat{\tau}_k)}{L_{\max}}$. We train π_θ with a PPO-style actor–critic objective [19, 17], since ground-truth target-direction labels are unavailable for the joint coverage-and-constraint objective.

4 Experiments

4.1 Dataset and setup

We evaluated on 8 retrospective temporal lobe epilepsy cases (84 electrodes) with CT/MRI-derived segmentations. Training/testing used patient-level 4/4 splits, so electrodes from the same patient were never used as independent train/test subjects; the 84 electrodes are implantation instances nested within patient maps. Although the number of patients is pilot-scale, each case requires detailed vascular/skull reconstruction and a faithful patient-specific planning environment. Each electrode first interacts with its own target and anatomical corridor, and then affects the full plan because accepted trajectories constrain later electrodes through spacing and collision rules. This makes the 84 electrode-implantation instances meaningful for evaluating the planner while preserving the patient-level split. ROIs included clinically common temporal targets such as left fusiform gyrus (L-FG), left inferior occipital gyrus (L-IOG), left hippocampus head (L-HPC-H), and right amygdala (R-Amyg). We compared against (i) manual plans from an experienced neurosurgeon (Doctor), (ii) CAP augmented with spatial priors (CAP+Priors) [8], and (iii) a recording-sensitivity optimizer (RS-Opt) [10]. All methods were evaluated under the same hard constraints in Eq. (2); our goal is plan quality under comparable time rather than runtime advantage. Implementation code will be released upon publication.

4.2 Evaluation metrics

We report capsule coverage c_k in Eq. (1), where higher coverage means more hypothesized epileptogenic gray matter lies within the electrode's effective sensing range. To quantify multi-electrode coordination, we measure the minimum inter-electrode spacing, defined as the minimum distance between any pair of planned trajectories in the same patient. Safety-related metrics include mean margins to vessels/deep structures, a CAP-style vessel-risk score, and the collision-free rate under Eq. (2). The margin summarizes SDF-based clearance to critical structures, while the risk score penalizes vessel-proximal paths; together with collision-free rate, these metrics assess evaluated safety under the same anatomical thresholds. We also report path length and average runtime per plan.

4.3 Results

Table 1 reports quantitative results on 8 retrospective cases (84 electrodes) under identical anatomical inputs and evaluation thresholds. The proposed method achieves the best ROI coverage rate (10.2%). Higher coverage means that more gray matter inside the hypothesized epileptogenic ROI falls within the electrode's effective sensing range, increasing the chance of capturing early ictal activity. It also yields the largest SDF-based margin (1.88 mm) and the lowest risk score (0.917). Compared with the clinical plan, ROI coverage increases from 5.6% to 10.2% while the margin improves from 1.08 mm to 1.88 mm, indicating stronger

Table 1. Quantitative comparison on 8 retrospective cases (84 electrodes). CAP (spatial priors) and RS-Opt are re-implemented baselines [8, 10]. Values are averaged over electrodes; spacing is plan-level min/avg (mm), time is minutes/case, and Collision-Free counts electrodes satisfying Eq. (2). Best values are bold.

Method	ROI Cov.↑	Safety Margin↑	Risk Score↓	Path Length↓	Coll.-Free↑	Spacing↑ (min/avg)	Time (min)↓
Clinical (Expert)	5.6%	1.08	1.055	74.3	**76/84 (90.5%)**	**10.2**/12.7	90
CAP (spatial priors) [8]	7.0%	1.00	1.024	54.4	50/84 (59.5%)	4.9/8.0	15
RS-Opt [10]	6.7%	1.20	1.107	60.5	55/84 (65.5%)	6.8/**13.6**	**14**
Proposed	**10.2%**	**1.88**	**0.917**	**52.2**	74/84 (88.1%)	9.2/10.9	28

functional coverage with preserved evaluated safety margins. The planned paths are also shorter on average (52.2 mm vs. 74.3 mm for the clinical plan), which is clinically desirable because shorter intracerebral trajectories reduce the amount of tissue traversed and can simplify implantation while still reaching the intended ROI.

Although the proposed method does not attain the highest collision-free rate (74/84, 88.1%), it is close to the clinical plan (76/84, 90.5%), with only two fewer electrodes satisfying all constraints, and outperforms CAP and RS-Opt (59.5% and 65.5%). The minimum inter-electrode spacing is likewise comparable to the clinical plan (9.2 mm vs. 10.2 mm), supporting coordination under the spacing constraint in Sec. 3.2. The method reduces planning time from 90 to 28 minutes per case; although slower than automated baselines (14–15 min), it provides stronger overall plan quality under the same protocol.

Target localization results. Fig. 2 visualizes capsule-to-ROI overlap for a representative electrode targeting the left fusiform gyrus. Under identical trajectory constraints, the proposed method refines the centroid initialization to better align the electrode capsule with the ROI, achieving 7.0% coverage vs. 3.5% for the clinical plan and 0.8%/5.8% for CAP+Priors/RS-Opt [8, 10].

Multi-electrode cooperative planning and constraint satisfaction. Fig. 3 compares full bundles in a representative 11-electrode case.

CAP+Priors [8] often produces unnecessarily long paths (L-IApex, R-Amyg). RS-Opt [10] may generate clinically implausible entries, e.g., a contralateral trajectory to the left L-HPC-H target and an occipital entry for R-Amyg, increasing collision risk. Our method yields ipsilateral bundles with reasonable entry sites while satisfying the spacing constraint (Sec. 3.2).

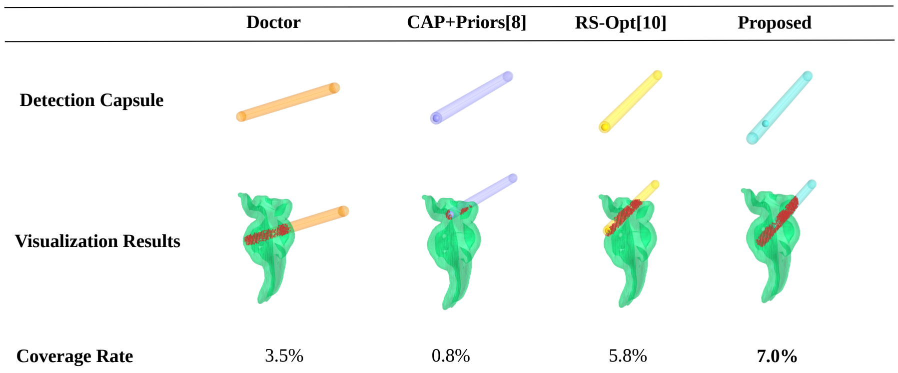

Fig. 2. Capsule-to-ROI overlap on the left fusiform gyrus. The proposed method achieves higher ROI coverage (7.0%) than the clinical plan (3.5%) and CAP/RS-Opt (0.8%, 5.8%) [8, 10].

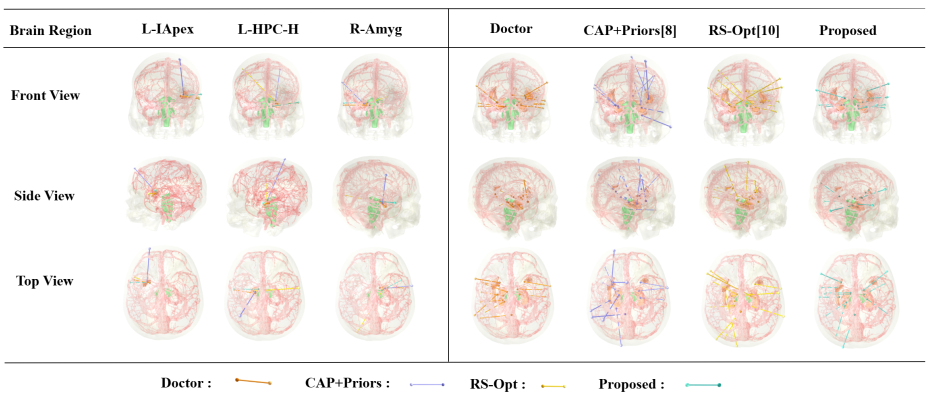

Fig. 3. Qualitative trajectory comparison. Left: three representative target regions (L-IApex, L-HPC-H, and R-Amyg) highlighting common failure modes. Right: a representative 11-electrode case comparing Doctor, CAP+Priors [8], RS-Opt [10], and the proposed method. CAP+Priors may produce unnecessarily long implantation paths, while RS-Opt can yield clinically implausible entry locations or hemisphere-crossing trajectories. The proposed method produces shorter, clinically plausible, and spatially coordinated trajectories under the same evaluation protocol.

5 Conclusion

We present a reinforcement-learning-based framework for multi-electrode SEEG planning that jointly localizes targets and generates straight-line trajectories under predefined anatomical and geometric constraints. By modeling an electrode's effective sensing range as a capsule, the proposed method directly optimizes ROI coverage during target localization, mitigating the decoupling between target selection and trajectory search common in prior pipelines. Deterministic projection and SDF-based constraint checking support fixed validation and cooperative multi-electrode planning via explicit inter-electrode spacing constraints. Experiments on 8 retrospective epilepsy cases (84 electrodes) support the promise of the coverage-aware planning approach, especially for improving ROI coverage while maintaining evaluated safety margins and coordinated inter-electrode spacing.

However, our work has several limitations, including the pilot-scale dataset, reliance on pre-segmented anatomical structures, and the use of straight-line trajectories, which may not fully reflect clinical preferences in highly constrained scenarios. Larger multi-center validation is needed to test robustness across heterogeneous implantation conventions, imaging protocols, and target hypotheses. In future work, we will investigate uncertainty-aware anatomical segmentation and explicit vascular modeling to improve robustness, and incorporate clinical priors and clinician-in-the-loop refinement to enhance routine clinical usability.

Acknowledgments. This work was supported in part by the National Natural Science Foundation of China (Nos. 62372441, 62303442), the Guangdong Basic and Applied Basic Research Foundation (Nos. 2026B1515020093, 2025A1515010423, 2024A1515140078), the Shenzhen Science and Technology Program (Grant No. RCYX20231211090127030), and the Research Program of the Department of Education of Guangdong Province (Grant No. 2025ZDZX3027).

Disclosure of Interests. The authors declare no competing interests.

References

1. Perucca, E., Perucca, P., White, H.S., Wirrell, E.C.: Drug Resistance in Epilepsy. Lancet Neurol. 22(8), 723–734 (2023).
2. Duncan, J.S., Taylor, P.N.: Optimising Epilepsy Surgery. Lancet Neurol. 22(5), 373–374 (2023).
3. Delgado-Garcia, G., Frauscher, B.: Future of Neurology & Technology: Stereoelectroencephalography in Presurgical Epilepsy Evaluation. Neurology 98(4), e437–e440 (2022).
4. Wu, S., Issa, N.P., Rose, S.L., et al.: Depth versus Surface: A Critical Review of Subdural and Depth Electrodes in Intracranial Electroencephalographic Studies. Epilepsia 65(7), 1868–1878 (2024).
5. Szmidel, M., Hunn, M., Neal, A., et al.: Vascular Imaging for Stereoelectroencephalography: A Safety and Planning Study. J. Clin. Neurosci. 127, 110762 (2024).
6. Liakina, T., Bartley, A., Carstam, L., Rydenhag, B., Nilsson, D.: Stereoelectroencephalography for Drug Resistant Epilepsy: Precision and Complications in Stepwise Improvement of Frameless Implantation. Acta Neurochir. 167(1), 75 (2025).
7. Girgis, F., Royz, E., Kennedy, J., Seyal, M., Shahlaie, K., Saez, I.: Superior Accuracy and Precision of SEEG Electrode Insertion with Frame-Based vs. Frameless Stereotaxy Methods. Acta Neurochir. 162(10), 2527–2532 (2020).
8. Dasgupta, D., Elliott, C.A., O'Keeffe, A.G., et al.: Computer-Assisted Stereoelectroencephalography Planning: Center-Specific Priors Enhance Planning. Front. Neurol. 16, 1514442 (2025).
9. Spiegler, P., Abdelsalam, H., Hellum, O., Hadjinicolaou, A., Weil, A.G., Xiao, Y.: PreVISE: An Efficient Virtual Reality System for SEEG Surgical Planning. Virtual Real. 29(1), 13 (2025).

10. Dessert, G.E., Thio, B.J., Grill, W.M.: Optimization of patient-specific stereo-EEG recording sensitivity. Brain Commun. 5(6), fcad304 (2023).
11. Wang, H.E., Woodman, M., Triebkorn, P., et al.: Delineating epileptogenic networks using brain imaging data and personalized modeling in drug-resistant epilepsy. Sci. Transl. Med. 15, eabp8982 (2023).
12. Thomas, J., Abdallah, C., Cai, Z., et al.: Investigating Current Clinical Opinions in Stereoelectroencephalography-Informed Epilepsy Surgery. Epilepsia 65(9), 2662–2672 (2024).
13. Chen, Y.W., Hanak, B.W., Yang, T.C., Wilson, T.A., Hsia, J.M., Walsh, H.E., et al.: Computer-assisted surgery in medical and dental applications. Expert Rev. Med. Devices 18(7), 669–696 (2021).
14. Sparks, R., Zombori, G., Rodionov, R., et al.: Automated Multiple Trajectory Planning Algorithm for the Placement of Stereo-Electroencephalography (SEEG) Electrodes in Epilepsy Treatment. Int. J. Comput. Assist. Radiol. Surg. 12(1), 123–136 (2017).
15. Li, R., Shi, Y., Si, W., et al.: Versatile Multi-Constrained Planning for Thermal Ablation of Large Liver Tumors. Comput. Med. Imaging Graph. 94, 101993 (2021).
16. Ao, Y., Esfandiari, H., Carrillo, F., et al.: SafeRPlan: Safe Deep Reinforcement Learning for Intraoperative Planning of Pedicle Screw Placement. Med. Image Anal. 99, 103345 (2025).
17. Amadou, A.A., Singh, V., Ghesu, F.C., et al.: Goal-Conditioned Reinforcement Learning for Ultrasound Navigation Guidance. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. Springer Nature Switzerland (2024).
18. Johari, M.M., Carta, C., Fleuret, F.: ESLAM: Efficient Dense SLAM System Based on Hybrid Representation of Signed Distance Fields. In: Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), pp. 17408–17419 (2023).
19. Son, S., Zheng, L., Sullivan, R., Qiao, Y.-L., Lin, M.: Gradient Informed Proximal Policy Optimization. In: Advances in Neural Information Processing Systems, vol. 36. Curran Associates, Inc. (2023).