Differentiable Manifold Optimization for Fast, Multi-Needle Ablation Planning

Philip Flörl[1], Matthias Harders[1], Reto Bale[2(✉)], and Stefano Fogarollo[1]

[1] Department of Computer Science Interactive Graphics and Simulation Group (IGS), University of Innsbruck, Technikerstraße 21, Innsbruck, Austria
[2] Interventional Oncology, Stereotaxy & Robotics (SIP), Department of Radiology, Medical University Innsbruck, Innsbruck, Austria
Reto.Bale@i-med.ac.at

Abstract. Stereotactic radiofrequency ablation is a minimally invasive, image-guided intervention that can treat large tumors using multiple coaxial ablation needles or applicators. Planning such interventions under strict clinical constraints is time-consuming and highly dependent on operator expertise. Existing automated approaches based on genetic algorithms, set-covering, or predefined heuristics have long runtimes and limited clinical applicability. In this work, we propose a novel differentiable planning framework that jointly optimizes many applicators, each described by a position in 3D space and an orientation given as a unit direction vector. We propose optimizing for sufficient coverage and safety of the trajectories using differentiable loss terms that also account for multiple pullback positions. On the public 3D-IRCADb-01 dataset, our method significantly improves the minimum ablative margin over related work in under 10 seconds (median), more than an order of magnitude faster. Further, we evaluated on a large multi-needle cohort spanning 176 lesions from 129 interventions at the Innsbruck University Hospital: the proposed plans are safe (distance from critical structures of 5.8 mm), effective (minimum ablative margin of 6.0 mm), feasible (minimum applicator inter-distance of 7.5 mm) and are similar to what the clinicians executed. To our knowledge, this is the first interventional ablation planning method to cast multi-needle trajectory planning as a differentiable, manifold-constrained problem and one of the first validated on a large-scale clinical cohort, with a runtime suitable for intraoperative use.

Keywords: differentiable optimization · treatment planning · thermal ablation

1 Introduction

Thermal ablation has emerged as a well-established alternative to surgical resection for the treatment of liver tumors. The randomized controlled COLLISION trial proved the non-inferiority of thermal ablation to resection for small colorectal liver metastases, with shorter hospital stays and fewer complications [15]. Radiofrequency ablation (RFA), the most widely adopted variant, is a minimally

invasive, image-guided procedure in which alternating current delivered through the applicator electrodes induces localized coagulative necrosis, thus achieving targeted tumor destruction while sparing surrounding healthy tissue [5]. Because a single electrode produces a limited coagulation volume, RFA has historically been used only for small lesions. To address this limitation, stereotactic radiofrequency ablation (SRFA) approaches have been developed [6]. Typically, contrast-enhanced CT images are obtained and used to plan the applicator trajectories. Then, multiple coaxial applicators are sequentially advanced to the predefined location, and a non-enhanced CT control scan is used to verify needle placement. Immediately after ablation, a scan is performed to assess ablation outcome [4,5].

Multi-needle planning, the focus of this paper, is a time-consuming task and requires specialized knowledge [2,24]: for each needle, clinicians must determine the target and skin insertion points such that the induced ablation zone covers the tumor alongside a safety margin, avoids critical structures such as bone, vasculature, and adjacent organs, and remains far apart from the other needles to prevent collisions. Early methods optimizing single-applicator plans include simplex optimization [3], GPU-accelerated Pareto-optimal trajectory selection [26], heat-distribution simulation [25], and reinforcement learning [9]. To target multi-applicator interventions, Chen et al. propose a clustering algorithm to automatically determine the optimal ablation zones [10]. Li et al. reduce the problem to 2D and optimize up to three applicators, with an average runtime of 59 s [16]. Wang et al. parameterize up to four cryoablation regions as 3D ellipsoids with differentiable Gaussians, with reported computation times ranging from minutes to hours [29]. In contrast, we optimize the applicator trajectories under differentiable safety and feasibility constraints. A further layer of complexity arises from the pullback technique, in which clinicians retract the active electrode along the trajectory to generate multiple contiguous ablation zones [4]. Liang et al. adopt a set-covering technique to optimize multiple needles with pullback with an average computation time of 8 min [18]. Yu et al. employ a genetic algorithm to find optimal trajectories for selected candidate target points [31]. Lukes et al. propose a genetic algorithm capable of handling multiple applicators with pullback, with a computation time that can reach several minutes in complex cases [19]. Li et al. developed a set of constraints within a multi-objective optimization framework targeting small, early-stage lesions [17]. Despite their differences, these methods often cast planning as a combinatorial search problem, which scales poorly with the number of needles and incurs runtimes that preclude intraoperative use.

This raises a central question: **how can we generate safe and feasible multi-needle plans for thermal ablation interventions intraoperatively?** Different from traditional combinatorial search methods, we propose to cast thermal ablation planning as a differentiable, manifold-constrained optimization problem. Our empirical results show that the proposed method improves the safety, efficacy, and feasibility of stereotactic radiofrequency ablation interventions, while being more than an order of magnitude faster than related work. In summary, the main contributions of this work are threefold: 1)

We formulate multi-needle planning as continuous optimization on the product manifold $\mathbb{R}^3 \times S^2$. 2) We replace combinatorial search with novel differentiable loss objectives to satisfy distance constraints and sufficient ablative margins. 3) We conduct comprehensive experiments on the 3D-IRCADb-01 public dataset and retrospectively evaluate on 176 lesions from the Innsbruck University Hospital, one of the largest multi-needle cohorts to our knowledge. In the following sections, we introduce our method, report the experiments, and discuss the implications.

2 Methods

We represent a treatment plan $\mathbf{p}$ as a set of $N_{\text{applicators}} \in \mathbb{N}$ applicators, each defined by a target point $\mathbf{t}_i \in \mathbb{R}^3$ and an insertion direction as a unit vector on the sphere manifold $\mathbf{d}_i \in S^2$. This parameterization is necessary to ensure smooth gradient updates during optimization [1,8] and to avoid gimbal lock discontinuities [32]. We approximate each ablation zone as a cylinder, following common practice for the pullback technique [19].

2.1 Anatomy Representation

We implicitly encode each anatomical structure (bones, vasculature, other organs) as a signed distance field (SDF) $\phi(\mathbf{x}, \partial\Omega)$, which maps a query point $\mathbf{x} \in \mathbb{R}^3$ to its signed Euclidean distance from the structure boundary (negative values inside $\partial\Omega$ and positive values outside). We precompute the SDFs on a uniform voxel grid and evaluate arbitrary sample points via trilinear interpolation, preserving first-order differentiability [23]. For structures consisting of multiple mesh boundaries (e.g., when combining different organs in a single mesh), we compute the SDF by taking the pointwise minimum distance to each SDF. Figure 1 shows a slice of an SDF volume and the corresponding plan proposed by our algorithm for a random case.

2.2 Loss Objective

In the following, let $\text{LSE}(\mathbf{x}) = \log \sum_i e^{x_i}$ be the LogSumExp-function with a temperature of 1.0, and accordingly $\mu(\mathbf{x}) = -\text{LSE}(-\mathbf{x})$ be a smooth approximation to the minimum. We define the total loss $\mathcal{L}$ to be optimized as:

$$\mathcal{L} = \lambda_{\text{safety}} \cdot \mathcal{L}_{\text{safety}} + \lambda_{\text{applicators}} \cdot \mathcal{L}_{\text{applicators}} + \lambda_{\text{coverage}} \cdot \mathcal{L}_{\text{coverage}} \tag{1}$$

where λ_* are the weights of each term explained below.

Distance From Critical Structures We extract bones, vasculature, and other organ meshes with marching cubes from segmentation masks obtained with TotalSegmentator [29,30]. To penalize trajectories with insufficient distance (less

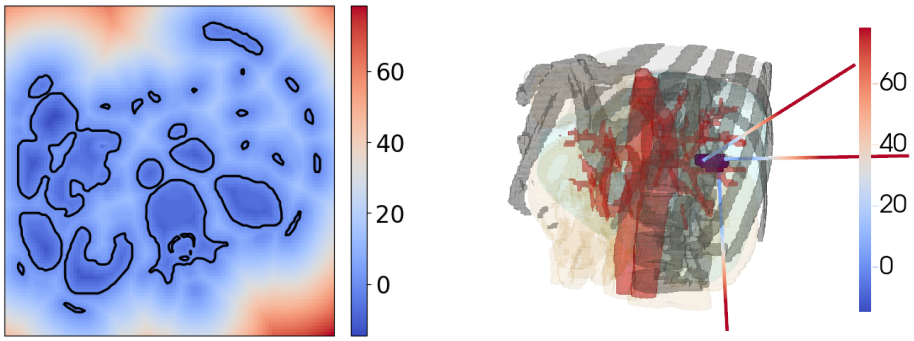

(a) Visualization of one slice of the SDF volume. Black contours show the outlines of anatomical structures.

(b) Plan generated by our algorithm with corresponding distances to nearest critical structures along each trajectory.

Fig. 1: We encode anatomical structures with a SDF grid. The colors show the signed distance in mm to the closest anatomical surface boundary: from red (far away) to blue (closer, possibly penetrating).

than α) from the boundary of critical structures $\partial\Omega_{\text{anatomy}}$, we define $\mathcal{L}_{\text{safety}}$ as:

$$\mathcal{L}_{\text{safety}}(\mathbf{p}) = \frac{1}{N_{\text{applicators}}} \sum_{i=1}^{N_{\text{applicators}}} \text{ReLU}\left(\alpha - \mu\left(\phi\left(\mathbf{s}_i, \partial\Omega_{\text{anatomy}}\right)\right)\right), \quad (2)$$

where $\mathbf{s}_i$ is a set of 1000 equidistant points spanning the probe length of 200 mm, sampled along the applicator's trajectory.

Applicator Inter-distance Let $\mathcal{S}$ be the segment-to-segment distance [11], then we define $\mathcal{L}_{\text{applicators}}$ to avoid applicator collisions as:

$$\mathcal{L}_{\text{applicators}}(\mathbf{p}) = \text{ReLU}\left(\beta - \mu_{i\neq j}(\mathcal{S}(\mathbf{p}_i, \mathbf{p}_j))\right). \quad (3)$$

Coverage Let $\mathbf{P}_{\text{tumor}} \in \mathbb{R}^{N_{\text{tumor}}\times 3}$ be the set of N_{tumor} tumor points to ablate and $\mathbf{P}_{\text{pullback}} \in \mathbb{R}^{N_{\text{applicators}}\times N_{\text{pullback}}\times 3}$ denote all pullback positions across all applicators. $\mathbf{P}_{\text{pullback},j,k} \in \mathbb{R}^3$ is the k-th pullback position along the j-th applicator, and N_{pullback} is the number of treatment points per applicator. In our experiments, we followed the protocol used by the clinicians at the Innsbruck University Hospital and used $N_{\text{pullback}} = 4$ evenly spaced at 30 mm intervals. We define the coverage loss as follows:

$$\mathcal{L}_{\text{coverage}}(\mathbf{P}_{\text{pullback}}, \mathbf{P}_{\text{tumor}}) = \frac{1}{N_{\text{tumor}}} \sum_{i=1}^{N_{\text{tumor}}} \mu_{j,k} \|\mathbf{P}_{\text{tumor},i} - \mathbf{P}_{\text{pullback},j,k}\|. \quad (4)$$

2.3 Iterative Optimization

The optimization workflow consists of the following steps:

1. for the initial applicator, we sample a random point from the tumor mesh as the target point. We set the initial direction vector parallel to the z axis, without loss of generality;
2. we use Riemannian Adam [8,13] to minimize Equation 1;
3. applicators are incrementally added and jointly optimized with all current ones until a target minimum ablative margin (MAM) is achieved, as the MAM is the clinically decisive quantity [14]. We initialize the target point of the new applicator to the tumor point farthest from the currently placed applicators and set the direction vector parallel to the z axis.

The target ablative margin is clinician-defined and tumor type dependent; following prior work, we used a margin of 5 mm in our experiments [19]. Heuristically, we progressively reduce the target ablative margin when >20 applicators are required and terminate optimization when >32 are needed.

We randomly sample a case and plot the progression of $\mathcal{L}$ and ablative margin in Figure 2a and the result in Figure 2b. In all our experiments, we optimize $\mathcal{L}$ for at most 1000 iterations with a patience of 100 iterations and a learning rate of 10^{-2}. For all datasets, we empirically set $\lambda_{\text{coverage}} = \lambda_{\text{safety}} = 1.0$, $\lambda_{\text{applicators}} = 0.1$, $\alpha = 2.0$ and $\beta = 7.0$. We implemented our algorithm in PyTorch and ran the experiments on Ubuntu 24.04 with an Intel Core i9-14900K CPU.

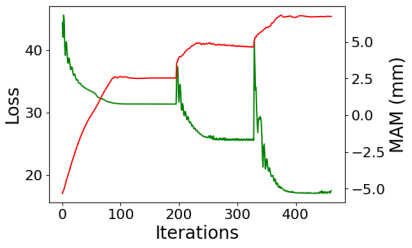

(a) Loss $\mathcal{L}$ (green) and MAM (red) progression on a case from the Innsbruck University Hospital. The insertion of a new applicator (at iterations 196 and 332) creates a spike in $\mathcal{L}$ due to collisions with anatomical structures and other applicators.

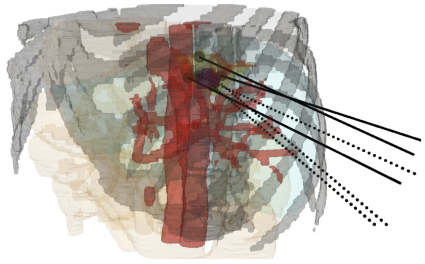

(b) Applicators inserted by clinicians (solid lines) vs. those proposed by our method (dashed lines). Liver tissue is shown in cyan, tumor in purple, vasculature in red, bones in gray, and other organs in orange.

Fig. 2: Loss and MAM progression (left) and corresponding plan proposed by our method with clinician's plan (right) as reference. In this case, three applicators are found to be sufficient to achieve a MAM of 6.8 mm, with a distance to the nearest critical structure of 4.9 mm.

3 Experiments & Results

3D-IRCADb-01 Cohort We follow related work and evaluate on the open-source 3D-IRCADb-01 dataset [19,28], which comprises 120 lesions from 20 patients; cases 5, 7, 11, 12, 14, and 20 were excluded due to the absence of liver tumors and an exophytic lesion in case 12 [19]. To compare with [19], we used their open-source code, while for [31,18], we relied on the metrics reported in their publications, as the respective authors did not share their code. Moreover, we filter our results by maximum tumor volume (TV) since [31,18] do not disclose the specific tumors used in their evaluations. For a fair comparison, we run our method with the smallest ablation zone radius (AZR).

We report the salient Mean Coverage (MC) and the 10% Trimmed Mean Coverage (TMC) metrics with an ablation margin of 5 mm. In addition, we present the MAM metric in mm as it is an independent predictor of local tumor progression, using 5 mm as a threshold for subcapsular and perivascular lesions [14]. In Table 1, we also show the number of lesions in the cohort (N), the number of applicators $N_{\text{applicators}}$ used for each method and the runtime T in seconds measured from loading the anatomical meshes, excluding the time spent on metric computation.

For all metrics, we report the median and the 5% and 95% confidence intervals in square brackets, with the arrows indicating whether higher (↑) or lower (↓) values are better. We denote statistically significant differences with the baseline (marked in italics) with an asterisk (*); to compute the p-value, we used the paired Wilcoxon signed-rank test with Bonferroni correction.

Table 1: Results on the 3D-IRCADb-01 dataset, grouped by tumor volume. A direct paired statistical comparison with [31] and [18] could not be performed, as neither study reports outcomes at the per-case level.

Method	AZR	N	TV	MC (TMC) ↑	MAM ↑	$N_{\text{applicators}}$ ↓	T ↓
ours	15	116	58.5	1.00 (1.00)	7.3 [5.0, 12.3]	1.0 [1.0, 7.2]	4.0 [2.5, 20.0]
[31]	15-40	32	44.5	0.99 (1.00)	N/A	1.6 [1.0, 3.5]	N/A
ours	23	115	27.6	1.00 (1.00)	14.6 [5.0, 20.3]	1.0 [1.0, 2.0]	3.5 [2.4, 6.1]
[18]	23-55	20	28.4	1.00 (1.00)	N/A	1.0 [1.0, 1.0]	480.0 [N/A]
ours	10	120	341.2	0.96 (0.99)	5.3 [2.2, 7.3]	4.0 [1.0, 22.0]	9.6 [2.5, 112.6]
[19]	10	120	341.2	0.80 (0.82)*	3.7 [0.3, 7.1]*	3.0 [1.0, 19.0]*	71.6 [3.4, 1689.9]*

Innsbruck University Hospital Cohort To assess clinical applicability, we performed a retrospective evaluation on 176 lesions across 129 SRFA interventions in 125 patients treated at the Innsbruck University Hospital. To our knowledge, this is one of the largest multi-needle cohort reported in the literature. First, we extracted the applicators from the intraoperative CT needle control scan by filtering based on Hounsfield intensities and organ segmentation masks

from TotalSegmentator [30]; then we performed line detection with RANSAC-3D [20]. Second, we non-rigidly registered this scan with the pre-interventional contrast-enhanced CT scan [27], in which two expert clinicians manually segmented the lesions; this co-registration is necessary to account for intraoperative deformations. For this cohort, we adopted an AZR of 15 mm, since insufficient ablative margins computed using a cylindrical 15 mm simulated ablation zone around the detected applicators were significantly associated with local tumor progression (p-value = 0.042).

We present results in Table 2, including the minimum distance to the closest anatomical critical structure (CS) and the minimum applicator inter-distance (ID), both in mm. In addition, we report a median and the 5% and 95% confidence intervals of the length of applicators inside the liver (↓) of 48.8 [11.1, 117.8] mm (compared to 60.0 [29.9, 100.1] and 38.8 [5.3, 105.5] mm obtained by clinicians and [19], respectively), and an intersection angle to liver surface (↑) of 37.6 [15.3, 77.9] (compared to 35.0 [11.8, 68.1] and 31.4 [3.4, 70.2] mm obtained by clinicians and [19], respectively). Note the MAM for clinicians' plans: this is due to the extraction of needles from the first intraoperative control CT, which may include suboptimal positions. To evaluate the plausibility of the plans, we trained a SetVAE [12] to learn a compact representation of the clinician's plans; after training for 500 epochs, the held-out RMSE reconstruction error was 0.47 [0.36, 0.60]. We quantify the similarity of a given plan to clinical practice as its Euclidean distance from the clinicians' plans in this learned feature space, with smaller values indicating greater similarity; we plot this feature space in Figure 3a. Our plans are significantly closer to the clinicians' (0.52 [0.21, 1.07]) than those of [19] (0.97 [0.46, 1.69]*), indicating that our method produces trajectories more consistent with expert practice.

Table 2: Quantitative results on the dataset from the Innsbruck University Hospital.

Method	MAM ↑	CS ↑	ID ↑	$N_{\text{applicators}}$ ↓	T ↓
ours	6.0 [5.1, 10.2]	5.8 [−0.3, 10.5]	7.5 [2.7, 11.2]	2.0 [1.0, 11.0]	8.6 [4.5, 67.0]
clinicians	0.3 [−13.5, 6.9]*	2.5 [−0.9, 8.7]*	8.7 [2.5, 20.3]*	3.0 [1.0, 7.0]	N/A
[19]	5.3 [2.3, 10.3]*	1.6 [0.0, 17.8]*	7.2 [3.3, 15.6]	1.0 [1.0, 11.6]	252.4 [12.0, 1001.4]*

4 Discussion and Conclusion

The empirical results reported in Table 1 and Table 2 show statistically significant improvements over related work. While our method needs one more applicator (median) than [19], it also achieves a significantly better MAM, necessary for the success of the intervention, and is more than an order of magnitude faster (see Figure 3b). Note that while we do not optimize for length of appli-

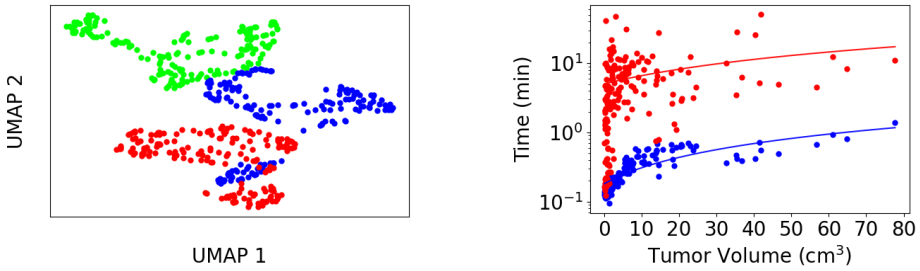

(a) UMAP projection [21] of the learned SetVAE feature space.

(b) Tumor volume and computation time (log scale); we fitted a line using least squares polynomial regression.

Fig. 3: On the left we visualize the feature space of our SetVAE while on the right we compare the tumor volume to the computation time. Our method (blue) is at the same time more similar to clinicians (green) and more than an order of magnitude faster than related work [19] (red).

Table 3: Ablation studies on the 3D-IRCADb-01 dataset, grouped by ablation area: parameterization of the applicators and loss terms.

Ablation	MAM ↑	CS ↑	ID ↑	$N_{\text{applicators}}$ ↓	T ↓
ours	5.3 [2.2, 7.3]	3.7 [−1.1, 6.9]	4.5 [0.2, 7.8]	4.0 [1.0, 22.0]	9.6 [2.5, 112.6]
[32]	5.3 [2.8, 7.0]	1.5 [−1.7, 6.0]*	4.3 [1.9, 7.4]	5.0 [1.0, 32.0]*	6.2 [2.0, 108.8]*
[7]	5.3 [3.0, 7.1]	0.0 [−2.9, 6.6]*	4.7 [2.9, 7.7]*	4.0 [1.0, 32.0]	11.8 [3.1, 175.4]*
$\mathcal{L}_{\text{safety}}$	5.2 [2.9, 7.1]	−2.1 [−7.9, 4.4]*	6.6 [3.6, 8.1]*	5.0 [1.0, 27.0]	10.9 [2.5, 114.7]
$\mathcal{L}_{\text{applicators}}$	5.3 [2.5, 7.3]	3.7 [−0.3, 7.3]	0.1 [0.0, 5.1]*	4.0 [1.0, 26.0]	9.0 [2.5, 125.3]
$\mathcal{L}_{\text{coverage}}$	3.3 [−2.2, 6.3]*	3.8 [−0.4, 6.8]	7.0 [0.1, 9.5]*	20.0 [1.0, 21.0]*	40.9 [2.6, 80.4]*

cator inside liver nor for the intersection angle to the liver surface, we achieve comparable results to related work and clinicians in these metrics.

Ablation Studies To evaluate the individual contribution of the proposed components, we conduct multiple ablation studies as shown in Table 3. First, we evaluate the impact of applicator representation by comparing our method with the Adam optimizer on the differentiable roto-translation manifold from [32]. We also compare our approach with Projected Gradient Descent (PGD) [7], parameterizing each applicator as a pair of insertion and target points: optimizing on the proposed manifold $\mathbb{R}^3 \times S^2$ yields significant improvements on safety distances and number of applicators over related representations [32], [7]. Finally, we run our method ablating each loss term individually (see the last three rows of Table 3). Including $\mathcal{L}_{\text{safety}}$ and $\mathcal{L}_{\text{applicators}}$ significantly improves the distance to critical structures and the applicator inter-distances, respectively. $\mathcal{L}_{\text{coverage}}$ helps to achieve a significantly higher MAM, with fewer applicators.

Limitations and Future Work We acknowledge several limitations of our current implementation. First, optimizing for applicator length in the liver and applicator angle to the liver surface could improve the feasibility of the trajectories. Second, the loss weights are fixed and independent of the number of applicators; our algorithm might benefit from an adaptive strategy. Third, differentiable models for ablation can be used to make a more realistic prediction of the MAM [22]. Finally, to initialize new applicators, future work could explore learning-based strategies. As a next step, we intend to address the limitations of our method and conduct multi-center studies evaluating the robustness and generalizability of our method. Although we focus on SRFA, our approach generalizes to microwave ablation, cryotherapy, and irreversible electroporation. In practice, the ablation-zone model, device constraints, and a subset of loss terms need to be adapted.

Conclusion In summary, we propose to cast multi-needle ablation planning as a differentiable, manifold-constrained optimization problem on $\mathbb{R}^3 \times S^2$. Our formulation can significantly improve the safety and feasibility at a fraction of the computational cost of related work. These results position the proposed framework as a viable step towards reliable, intraoperative planning for stereotactic thermal ablation interventions. The source code is available at `https://git.uibk.ac.at/informatik/igs/open/r-adam-srfa`.

Acknowledgements This project has been supported by the Austrian Science Foundation (10.55776/DOC110) and by the Österreichische Forschungsförderungsgesellschaft (FO999925705).

Disclosure of Interests The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Absil, P.A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press (2008)
2. Altrogge, I., Kröger, T., Preusser, T., Büskens, C., Pereira, P.L., Schmidt, D., et al.: Towards optimization of probe placement for radio-frequency ablation. In: MICCAI. pp. 486–493 (2006)
3. Baegert, C., Villard, C., Schreck, P., Soler, L., Gangi, A.: Trajectory optimization for the planning of percutaneous radiofrequency ablation of hepatic tumors. Computer Aided Surgery **12**(2), 82–90 (2007)
4. Bale, R., Schullian, P., Eberle, G., Putzer, D., Zoller, H., Schneeberger, S., Manzl, C., et al.: Stereotactic radiofrequency ablation of hepatocellular carcinoma: A histopathological study in explanted livers. Hepatology **69**(3), 1030–1040 (2019)
5. Bale, R., Widmann, G., Haidu, M.: Stereotactic Radiofrequency Ablation. CardioVascular and Interventional Radiology **34**(4), 852–856 (2011)
6. Bale, R., Widmann, G., Stoffner, D.R.: Stereotaxy: Breaking the limits of current radiofrequency ablation techniques. European Journal of Radiology **75**(1), 32–36 (2010)
7. Bolduc, E., Knee, G.C., Gauger, E.M., Leach, J.: Projected gradient descent algorithms for quantum state tomography. npj Quantum Information **3**(1), 44 (2017)
8. Bécigneul, G., Ganea, O.E.: Riemannian adaptive optimization methods. In: International Conference on Learning Representations (ICLR) (2019)
9. Chaitanya, K., Audigier, C., Balascuta, L.E., Mansi, T.: Automatic planning of liver tumor thermal ablation using deep reinforcement learning. In: Proceedings of The 5th International Conference on Medical Imaging with Deep Learning. pp. 219–230 (2022)
10. Chen, R., Jiang, T., Lu, F., Wang, K., Kong, D.: Semiautomatic radiofrequency ablation planning based on constrained clustering process for hepatic tumors. IEEE Transactions on Biomedical Engineering **65**(3), 645–657 (2018)
11. Ericson, C.: Real-Time Collision Detection. Morgan Kaufmann (2004)
12. Kim, J., Yoo, J., Lee, J., Hong, S.: SetVAE: Learning hierarchical composition for generative modeling of set-structured data. In: CVPR. pp. 15059–15068 (2021)
13. Kochurov, M., Karimov, R., Kozlukov, S.: Geoopt: Riemannian optimization in PyTorch (2020)
14. Laimer, G., Schullian, P., Jaschke, N., Putzer, D., Eberle, G., Alzaga, A., Odisio, B., Bale, R.: Minimal ablative margin (MAM) assessment with image fusion: an independent predictor for local tumor progression in hepatocellular carcinoma after stereotactic radiofrequency ablation. European Radiology **30**(5), 2463–2472 (2020)
15. van der Lei, S., Puijk, R.S., Dijkstra, M., Schulz, H.H., Vos, D.J.W., De Vries, J.J.J., et al.: Thermal ablation versus surgical resection of small-size colorectal liver metastases (COLLISION): an international, randomised, controlled, phase 3 non-inferiority trial. The Lancet Oncology **26**(2), 187–199 (2025)
16. Li, R., An, C., Wang, S., Wang, G., Zhao, L., Yu, Y., et al.: A heuristic method for rapid and automatic radiofrequency ablation planning of liver tumors. International Journal of Computer Assisted Radiology and Surgery **18**(12), 2213–2221 (2023)
17. Li, S., Zhou, F., Zhang, Y., Xu, S., Wang, Y., Cheng, L., Bie, Z., Li, B., Li, X.G.: Multi-stage automatic and rapid ablation and needle trajectory planning method for CT-guided percutaneous liver tumor ablation. Medical Physics **52**(1), 113–130 (2025)

18. Liang, L., Cool, D., Kakani, N., Wang, G., Ding, H., Fenster, A.: Automatic radiofrequency ablation planning for liver tumors with multiple constraints based on set covering. IEEE Transactions on Medical Imaging **39**(5), 1459–1471 (2020)
19. Lukes, A., Bale, R., Freysinger, W.: Automatic trajectory planning for stereotactic radiofrequency ablation in non-discrete search space. International Journal of Computer Assisted Radiology and Surgery **20**(7), 1349–1360 (2025)
20. Mariga, L.: pyRANSAC-3D (2022), `https://github.com/leomariga/pyRANSAC-3D`
21. McInnes, L., Healy, J., Saul, N., Großberger, L.: Umap: Uniform manifold approximation and projection. Journal of Open Source Software **3**(29), 861 (2018)
22. Mehtali, J., Verde, J.M., Essert, C.: C-NCA: Chained Neural Cellular Automata for Fast and Accurate Thermal Ablation Estimation . In: MICCAI. pp. 67 – 77 (2025)
23. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR. pp. 165–174 (2019)
24. Schullian, P., Laimer, G., Johnston, E., Putzer, D., Eberle, G., Widmann, G., Scharll, Y., Bale, R.: Reliability of stereotactic radiofrequency ablation (SRFA) for malignant liver tumors: Novice versus experienced operators. Biology **12**(2) (2023)
25. Schumann, C., Rieder, C., Haase, S., Teichert, K., Süss, P., Isfort, P., et al.: Interactive multi-criteria planning for radiofrequency ablation. International Journal of Computer Assisted Radiology and Surgery **10**(6), 879–889 (2015)
26. Seitel, A., Engel, M., Sommer, C.M., Radeleff, B.A., Essert-Villard, C., Baegert, C., et al.: Computer-assisted trajectory planning for percutaneous needle insertions. Medical Physics **38**(6Part1), 3246–3259 (2011)
27. Siebert, H., Großbröhmer, C., Hansen, L., Heinrich, M.P.: ConvexAdam: Self-Configuring Dual-Optimization-Based 3D Multitask Medical Image Registration. IEEE Transactions on Medical Imaging **44**(2), 738–748 (2025)
28. Soler, L., Hostettler, A., Agnus, V., Charnoz, A., Fasquel, J., Moreau, J., et al.: 3D image reconstruction for comparison of algorithm database: A patient specific anatomical and medical image database. Tech. rep., IRCAD, Strasbourg, France (2010)
29. Wang, C., Huang, X., Kong, Y., Li, Q., Hao, Y., Zhou, X.: 3DGPS: A 3D differentiable-gaussian-based planning strategy for liver tumor cryoablation. In: MICCAI 2024. pp. 3–13 (2024)
30. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., et al.: TotalSegmentator: Robust segmentation of 104 anatomic structures in CT images. Radiology: Artificial Intelligence **5**(5), e230024 (2023)
31. Yu, P., Fu, T., Wu, C., Jiang, Y., Yang, J.: Automatic radiofrequency ablation planning for liver tumors: A planning method based on the genetic algorithm with multiple constraints. In: Proceedings of the 2021 International Conference on Intelligent Medicine and Health. pp. 8–14 (2022)
32. Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: CVPR. pp. 5738–5746 (2019)