Learning Temperature-Dependent Properties in Cryoablation with Differentiable Physics

Raunak Samanta[1,2], Adarsh Raghunath[*,1,2], Chloé Audigier[3], Teodor Ionut Matei[4], Oliver Frings[2], Andreas Maier[1], and Felix Meister[2]

[1] Pattern Recognition Lab, Department of Computer Science, Friedrich-Alexander University Erlangen-Nürnberg, Erlangen, Germany
[2] Siemens Healthineers, Digital Technology and Innovation, Erlangen, Germany
[3] Siemens Healthineers, Swiss Innovation Hub, Lausanne, Switzerland
[4] Siemens S.R.L., Foundational Technologies, Brașov, Romania

*Corresponding author:adarsh.raghunath@fau.de

Abstract. Cryoablation is a minimally invasive tumor treatment whose success depends on accurate procedure planning under patient-specific anatomical and biophysical variability. Computational models could facilitate planning by forecasting post-operative ablation zones. However, they require temperature-dependent tissue properties for highly accurate predictions, which are difficult to measure across cryogenic ranges. These properties could instead be inferred by personalizing such models to experimental data, yet phase-change-induced nonlinearities often require flexible function representations with potentially many tunable parameters, making conventional optimization methods loose tractability. We propose a differentiable physics framework to personalize a cryoablation bioheat model from temperature observations using gradient-based optimization. Tissue properties are represented as bounded B-splines and optimized end-to-end, with a curvature penalty to reduce non-physical oscillations. In synthetic bench test experiments, we evaluated the method on target functions of increasing complexity and compared it with BFGS and Nelder-Mead. Differentiable-physics-based optimization remained effective as the number of trainable parameters increased, achieving competitive or better recovery of tissue properties, while requiring substantially fewer function evaluations than the comparators. These results suggest that differentiable physics is a parametrically scalable approach for estimating non-linear thermal tissue properties in cryoablation modeling.

Keywords: Cryoablation · Computational Modeling · Differentiable Physics · Digital Twin · Parameter Estimation.

1 Introduction

Cryoablation is a minimally invasive thermal ablation modality used to treat selected tumors by freezing tissue to cytotoxic temperatures, thereby inducing cellular necrosis [9]. Successful treatment with a low chance of local tumor recurrence requires meticulous pre-procedural planning, which currently relies heavily on clinician expertise and imaging [3]. Manufacturers support this planning process by providing reference charts that summarize expected ablation performance under controlled bench test conditions. The underlying data is typically generated in conductive phantom media (e.g., ultrasound gel heated to body core temperature), using standardized freeze–thaw protocols (e.g., a 10-5-10 min freeze-thaw-freeze cycle) and thermocouple sensors [2, 7, 16]. Resulting temperature isocontours are translated into the reference charts used for procedural planning. While useful, they only capture certain pre-defined probe placements and fail to capture patient-specific variability or complex in-vivo biophysics. Computational modeling could be used instead to forecast the growth of the iceball and the resulting post-operative ablation zone for arbitrary probe configurations and surrounding tissues. Hence, it may serve as a valuable tool for treatment planning, reducing the cognitive burden on clinicians [13].

Most computational models simulate cryoablation using the Pennes bioheat equation and often approximate the thermal tissue properties as constants or piecewise functions across phase-change thresholds [5, 10, 13, 16]. However, predictive performance is often limited because these simplified assumptions do not capture the true phase-change behavior of the tissue, where thermal conductivity changes abruptly and effective volumetric heat capacity exhibits a sharp peak [2, 7]. However, accurately measuring tissue thermal properties across the full cryogenic range is technically demanding and requires specialized equipment.

An alternative route is data-driven personalization of the computational model using observational data from cryoablation experiments. Here, it may be necessary to use flexible, high-dimensional parameterizations rather than simple baseline functions to capture the underlying strongly non-linear behavior. Traditional computational model personalization methods (e.g., parameter sweeps [18], simplex [6], or other methods [1, 4]), however, require repeated, expensive forward evaluations, rapidly losing tractability as parameter dimensionality increases. Differentiable physics offers a promising alternative by enabling a parametrically scalable, gradient-based optimization of high-dimensional models. By computing exact end-to-end sensitivities via automatic differentiation across time-stepping dynamics, it enables highly effective inverse parameter estimation [8, 12].

Motivated by these advantages, we developed a differentiable implementation of the simplified Pennes bioheat model for cryoablation. The framework models thermal conductivity and volumetric heat capacity as trainable temperature-dependent functions and uses an objective function combining data fidelity with curvature-based regularization. In addition, we conducted a series of synthetic experiments to assess the scalability and practical effectiveness of the proposed framework.

2 Methodology

2.1 Simulating Cryoablation In A Bench Test Setup

In this proof-of-concept study, we aim to uncover tissue-specific thermal properties from spatiotemporal temperature measurements using a synthetic setup consistent with single-probe cryoablation bench testing. In particular, we model cryoablation in a heated ultrasound gel phantom as an inhomogeneous heat diffusion problem defined by

$$c_v(\mathbf{T})\frac{\partial T}{\partial t} = \nabla \cdot k(\mathbf{T})\nabla\mathbf{T} \tag{1}$$

where $\mathbf{T}$ is the temperature field, t is time, and k and c_v are the temperature-dependent thermal conductivity and volumetric heat capacity, respectively. To solve the equation numerically with the Finite-Difference method, our simulations were performed on a 2D computational domain of 120×120 mm (see Fig. 1), discretized with a uniform cell size of 2 mm and a time step of 0.2 s, chosen to always satisfy the CFL condition in our experiments while keeping computational cost and training time manageable.

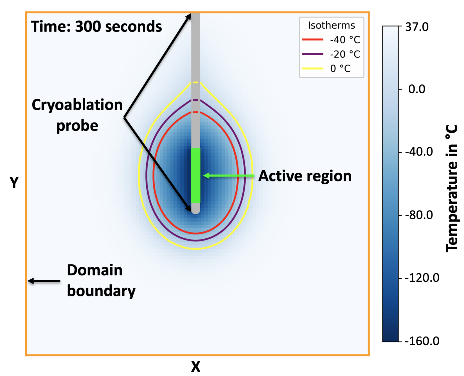

Fig. 1. Illustration of the 2D computational domain used to simulate heat diffusion during cryoablation, showing the temperature field at 300s into the freezing cycle.

The domain comprises two compartments: (1) the probe which was modeled as Stainless Steel 304[5] with a diameter of 2.1 mm and its non-insulated active zone of length 16 mm, located 3 mm away from the probe tip and positioned such that it is located in the domain center; and (2) the surrounding thermal conductive medium, initialized at 37 °C and modeled to behave thermally similar to ultrasound gel. In this work, we chose three configurations of varying complexity

[5] https://trc.nist.gov/cryogenics/materials/304Stainless/304Stainless_rev.htm

to model the thermal properties of the gel and to generate the synthetic reference data for our experiments. The first configuration describes a linear temperature dependence, which we model as

$$k_1(T) = (-0.0018 \cdot T + 0.08) \cdot 0.01, \quad \text{and} \quad c_{v1}(T) = (0.02 \cdot T + 4) \cdot 0.001. \quad (2)$$

The second configuration defines a simple non-linear temperature dependence, which is a handcrafted and smoother variant of the mathematical equations provided in [7], given by

$$k_2(T) = \begin{cases} (0.00012 \cdot T^2 + 2.3) \cdot 0.001, & \text{if } T < -5°C, \\ (-\tanh(T + 20) + 1.5) \cdot 0.001, & \text{if } T \geq -5°C, \end{cases} \quad (3)$$

$$c_{v2}(T) = \begin{cases} (-0.8 \cdot \log(-T) + 5.1) \cdot 0.001, & \text{if } T < -20°C, \\ (0.65 \cdot \tanh(0.2 \cdot T + 2) + 3.3) \cdot 0.001, & \text{if } T \geq -20°C. \end{cases} \quad (4)$$

The final configuration[6] was manually constructed from the sparse thermal property measurements provided in [7], which suggests a significant spike in the volumetric heat capacity around the freezing point. To this end, temperature intervals of distinct characteristic were first defined and plausible reference functions were manually fitted to these distinct intervals.

Finally, we apply a Dirichlet boundary condition to the domain boundary, keeping the temperature constant at 37 °C, and another Dirichlet boundary condition to the active zone of the probe, which progressively decreases the prescribed temperature from 21°C (room temperature) to -80°C within 30 s and then to -160 °C over the next 270 s linearly.

2.2 Learning Thermal Properties From Temperature Measurements

To uncover the thermal dependence of the tissue properties, we propose a differentiable-physics-based pipeline (see Fig. 2). Given an initial temperature field, the differentiable Finite-Difference solver with a set of tuneable parameters θ computes the temperature evolution of the bench test cryoablation for a predefined look-ahead horizon, selected to be equal to the largest time stamp within the reference data. Next, a loss function computes the mismatch between the simulated temperatures and the reference supervision points. Backpropagation is then applied to obtain gradients w.r.t. the underlying parameters. Since the backpropagation increases the computational cost significantly, we may further limit the amount of backpropgation steps by first running the simulation for a selectable warm up period without tracking the gradient information, before continuing the simulation with gradient tracking enabled. Finally, a gradient-descent-based optimizer is updating the parameters and the loop is repeated until a pre-defined threshold is reached or the optimization converged.

[6] The mathematical formulation requires 10 equations and exceeds the space constraint. The formulas can be obtained upon request.

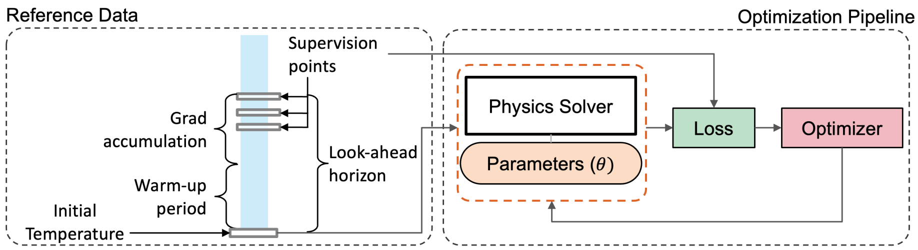

Fig. 2. Schematic overview of the training pipeline. Our differentiable physics solver predicts the temperature evolution from the initial state, with the underlying parameters (θ) iteratively updated by the Adam optimizer based on the loss comparing the reference temperature fields at discrete supervision points against the computed ones.

Trainable Representation Of Thermal Properties The proposed differentiable pipeline as described above would already enable the parameter optimization of existing mathematical material models, as long as these are differentiable. In this work, however, we chose instead a more flexible approach that does not assume prior knowledge about the functional form of the temperature dependence, except knowing appropriate bounds for the property range. Specifically, we model the conductivity and the volumetric heat capacity function each as a B-Spline with an adaptable total number of free parameters θ, initialized with Xavier initialization and divided equally between each B-Spline. Moreover, given the unconstrained B-Spline $\tilde{m}_p$ for a material property $p \in \{k, c_v\}$, we constrain the output to the pre-defined lower and upper limits subject to the transformation

$$m_p(T) = m_{p,\min} + (m_{p,\max} - m_{p,\min})\ \sigma(\tilde{m}_p(T))\,, \tag{5}$$

where $m_{p,\min}$ and $m_{p,\max}$ are the prescribed lower and the upper limit of the material property and $\sigma(\cdot)$ denotes the sigmoid function. This guaranties that $m_{p,\min} < m_p(T) < m_{p,\max}$ for all T. Lower bounds were imposed to ensure non-negativity and numerical stability, while upper bounds were set based on target ranges to maintain physical plausibility and prevent solver instability.

Loss Function The total loss is defined as $L = L_{primary} + w \cdot L_{smooth}$ with

$$\mathcal{L}_{primary} = \begin{cases} 0.5r^2, & \text{if } |r| \le 1, \\ |r| - 0.5, & \text{if } |r| > 1, \end{cases} \tag{6}$$

and where $r = y - \hat{y}$ is the prediction error between predicted and reference temperature fields at selected supervision points. To reduce non-physical oscillations, a smoothing loss $\mathcal{L}_{smooth} = \mathcal{R}_\delta(k) + \mathcal{R}_\delta(c_v)$, is added as a curvature penalty to regularize second-order variations [11]. For a given function $f \in \{k, c_v\}$ and a set of temperature samples $\mathcal{I}$, we define $\mathcal{R}_\delta(\cdot)$ as

$$\mathcal{R}_\delta(f) = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \left(f''(T_i)\right)^2\,, \tag{7}$$

and approximate $f''(T_i)$ using the central finite difference scheme

$$f''(T_i) \approx \frac{f(T_i+\delta) - 2f(T_i) + f(T_i-\delta)}{\delta^2} \tag{8}$$

with δ set to 10^{-2} and $\mathcal{I}$ comprising a set of equally spaced points over the temperature range of $-160°$C to $40°$C. Since L_{smooth} can vary over several orders of magnitude initially, the weighting factor w must be adjusted to account for scale differences between the two terms. From our experience, Xavier uniform initialization of larger B-Spline-based models yields higher L_{smooth}. Therefore, its weight was chosen such that initially L_{smooth} was at least four orders of magnitude smaller than $L_{primary}$, enabling it to act only after the primary loss decreased. Since all experiments used fewer than 200 trainable parameters, $w = 10^7$ was applied throughout.

3 Experiments and Results

3.1 Comparative Analysis of Computational Efficiency

Differentiable physics with Adam as the optimizer was compared against BFGS [1] (second-order, Finite-Difference gradients) and Nelder–Mead [15] (gradient free) to evaluate scalability with increasing model size. A tolerance of 10^{-6} was used for convergence across all methods to ensure consistency. All experiments used a look-ahead horizon of 300 seconds with supervision every 15 seconds. The number of trainable parameters was varied between 10 and 80 parameters. As illustrated in Fig. 3, for the linear and simple non-linear targets, differentiable physics achieved good approximations, particularly at lower temperatures, though some oscillations were observed. Differentiable physics achieved performance comparable to or better than BFGS, while both methods generally outperformed Nelder-Mead. This trend was particularly evident in the complex non-linear case, where differentiable physics and BFGS produced more stable results, whereas Nelder-Mead exhibited larger oscillations and degraded with increasing model size. It was also observed that more parameters were required to capture the complex nonlinearities. As model size increased, differentiable physics consistently required two orders of magnitude fewer function evaluations than Nelder-Mead and BFGS. Although each differentiable-physics evaluation took about five times longer than a single forward simulation without gradient computation (10.03 seconds), the method was still up to ten times faster overall than Nelder-Mead and BFGS.

3.2 Temperature Field Evolution Prediction Accuracy

Since low training error in an ill-posed inverse problem does not necessarily imply accurate prediction of unseen temperature dynamics, we further evaluated the ability of differentiable-physics-trained material models to predict the temperature evolution beyond the look-ahead horizon. Models with 4 to 160 parameters

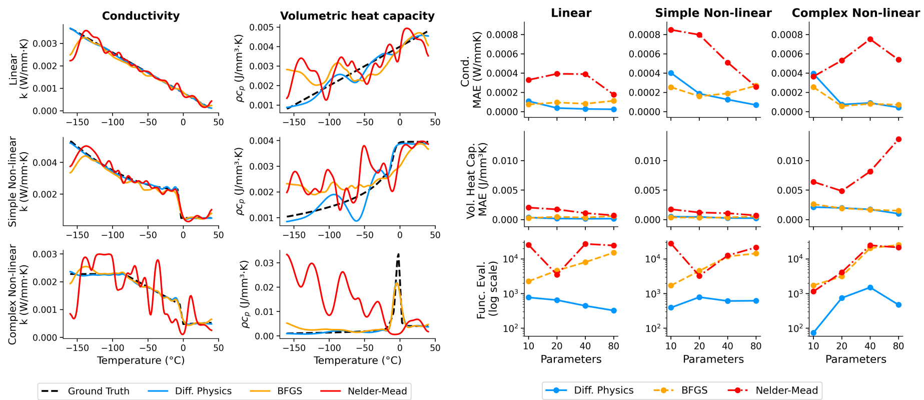

Fig. 3. (Left) Learned material property function representations from the model with 80 trainable parameters, overlaid on the three target functions of increasing complexity. (Right) Mean absolute error of the learned material property functions and the corresponding number of function evaluations, shown as a function of the number of model parameters for each target function.

were trained with the complex non-linear target data, a 300 s look-ahead horizon, and with either a single final supervision point (with 240 s warm-up period) or supervision every 15 s (no warm-up period). Performance was assessed by simulating a 10–5–10 min freeze–thaw–freeze cycle and computing the average MAE of the predicted temperature over the full 1500 s duration (see Fig. 4).

Models trained with multiple supervision points consistently outperformed those with a single point, except for the 4-parameter model, which lacks the capacity to capture complex non-linearities. Prediction errors at exact supervision points were lower than those across the complete 1500 s simulation. The

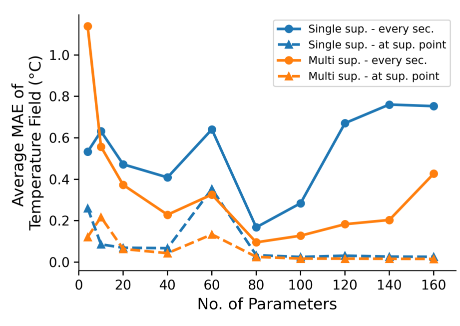

Fig. 4. Average mean absolute error (MAE) of predicted temperature fields as a function of the number of trainable parameters in the material model. Each model was evaluated by simulating a complete ablation protocol, with the MAE computed at every second and then averaged over the full simulation duration.

optimization explicitly minimizes errors at these markers, whereas the full simulation tests generalization. Notably, scaling beyond 80 parameters yielded only marginal gains at supervision points while degrading overall simulation accuracy. This degradation strongly indicates over-parameterization. A highly flexible B-spline learns a representation that can predict temperatures at the supervision points more accurately than a full freeze-thaw-freeze cycle.

4 Discussion and Conclusion

Our results demonstrate that differentiable physics provides a parametrically scalable and efficient framework for cryoablation model calibration and the discovery of temperature-dependent thermal properties. Compared to conventional optimization methods such as BFGS and Nelder-Mead, the proposed approach remains tractable as the number of parameters in the model increases.

The reported reduction in function evaluations also partly reflects that the comparators relied on finite-difference or gradient-free updates, whose cost scales with parameter dimensionality. A comparison against quasi-Newton methods using automatic-differentiation gradients (e.g., L-BFGS) remains an important direction for future work. Moreover, incorporating temporal information during training improves prediction quality, indicating that richer spatiotemporal data can help constrain the inverse problem.

We evaluated the feasibility of our proposed pipeline on idealized, noise-free synthetic data based on plausible handcrafted thermal property configurations. Future work will aim at transferring the pipeline to real bench test data, which includes noise and uncertainties from probe placement, thermocouple sensors, and gel impurities, among other sources. Furthermore, we used a simplified Pennes model without perfusion or metabolic heat, which is mainly appropriate for the current bench experiments. However, translation to patient-specific adaptive cryoablation models will require more clinically realistic data, such as isotherm boundaries obtained from intra-operative cryoablation imaging. A key challenge is that patient data may be sparse in time, which limits the amount of information available to constrain the inverse problem.

Moreover, estimating temperature-dependent thermal properties from temperature data remains inherently ill-posed. This is reflected in the oscillations observed in the learned functions, despite the low errors in the predicted temperature fields. It indicates that different material property functions, especially common scalings of k and c_v could result in the same thermal diffusivity (k / c_v ratio), which may explain the measured data equally well.

Future work should therefore further constrain the solution space by incorporating prior knowledge, such as physiologically motivated constraints or known functional behavior across temperature regimes published in literature, as explicit constraints or regularization terms. The observed oscillations may also indicate over-parameterization, particularly given that the inverse problem depends on both the chosen parameterization, and the spatial and temporal discretizations. Future work should evaluate the suitability of B-Splines for representing temperature-dependent material properties and compare their performance with alternative function approximation methods. In addition, since spatial grid resolution and time step size were kept constant across all experiments, a systematic investigation of spatial and temporal discretizations is therefore needed to determine an appropriate balance between accuracy, stability, and computational cost.

Reducing optimization time will also be important for time constrained clinical applications. This could be addressed through performance engineering, improved solver implementations, or recent physics-based deep learning approaches, such as the correction of discretization errors [17] and surrogate models [14].

The results demonstrate that our methodology enables computationally efficient personalization of cryoablation models by recovering temperature-dependent thermal properties from spatiotemporal temperature data. The findings establish both methodological feasibility and practical relevance of differentiable physics for high-fidelity cryoablation modeling.

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

Disclaimer. The concepts and information presented in this presentation are based on research results that are not commercially available. Future commercial availability cannot be guaranteed.

References

1. Quasi-Newton Methods. In: Nocedal, J., Wright, S.J. (eds.) Numerical Optimization, pp. 192–221. Springer-Verlag, New York (1999). https://doi.org/10.1007/0-387-22742-3_8, , series Title: Springer Series in Operations Research and Financial Engineering
2. Agafonkina, I.V., Belozerov, A.G., Berezovsky, Y.M., Korolev, I.A., Pushkarev, A.V., Tsiganov, D.I., Shakurov, A.V., Zherdev, A.A.: Thermal Properties of Biological Tissue Gel-Phantoms in a Wide Low-Temperature Range. Journal of Engineering Physics and Thermophysics **94**(3), 790–803 (May 2021). https://doi.org/10.1007/s10891-021-02356-z,
3. Ahmed, M., Brace, C.L., Lee, F.T., Goldberg, S.N.: Principles of and Advances in Percutaneous Ablation. Radiology **258**(2), 351–369 (Feb 2011). https://doi.org/10.1148/radiol.10081634,
4. Audigier, C., Mansi, T., Delingette, H., Rapaka, S., Passerini, T., Mihalef, V., Pop, R., Diana, M., Soler, L., Kamen, A., Comaniciu, D., Ayache, N.: Challenges to Validate Multi-physics Model of Liver Tumor Radiofrequency Ablation from Pre-clinical Data. p. 29 (Oct 2015),
5. Blezek, D.J., Carlson, D.G., Cheng, L.T., Christensen, J.A., Callstrom, M.R., Erickson, B.J.: Cell Accelerated Cryoablation Simulation. Computer Methods and Programs in Biomedicine **98**(3), 241–252 (Jun 2010). https://doi.org/10.1016/j.cmpb.2009.09.004,
6. Breen, M.S., Breen, M., Butts, K., Chen, L., Saidel, G.M., Wilson, D.L.: MRI-guided Thermal Ablation Therapy: Model and Parameter Estimates to Predict

Cell Death from MR Thermometry Images. Annals of Biomedical Engineering **35**(8), 1391–1403 (Aug 2007). https://doi.org/10.1007/s10439-007-9300-3,
7. Etheridge, M.L., Choi, J., Ramadhyani, S., Bischof, J.C.: Methods for Characterizing Convective Cryoprobe Heat Transfer in Ultrasound Gel Phantoms. Journal of Biomechanical Engineering **135**(2), 0210021–02100210 (Feb 2013). https://doi.org/10.1115/1.4023237,
8. Farsi, A., Bouziani, N., Ham, D.A.: Missing Physics Discovery through Fully Differentiable Finite Element-Based Machine Learning (Jul 2025). https://doi.org/10.48550/arXiv.2507.15787, , arXiv:2507.15787
9. Gage, A.A., Baust, J.: Mechanisms of Tissue Injury in Cryosurgery. Cryobiology **37**(3), 171–186 (Nov 1998). https://doi.org/10.1006/cryo.1998.2115,
10. Golkar, E., Rao, P.P., Joskowicz, L., Gangi, A., Essert, C.: GPU-based 3D iceball modeling for fast cryoablation simulation and planning. International Journal of Computer Assisted Radiology and Surgery **14**(9), 1577–1588 (Sep 2019). https://doi.org/10.1007/s11548-019-02051-8,
11. Green, P.J., Silverman, B.W.: Nonparametric regression and generalized linear models: a roughness penalty approach. Chapman and Hall, London (1994)
12. Karnakov, P., Litvinov, S., Koumoutsakos, P.: Solving inverse problems in physics by optimizing a discrete loss: Fast and accurate learning without neural networks. PNAS Nexus **3**(1), pgae005 (Dec 2023). https://doi.org/10.1093/pnasnexus/pgae005,
13. Li, S., Zhang, Y., Zhou, F., Xu, S., Wang, Y., Cheng, L., Bie, Z., Li, B., Li, X.G.: Iceball growth 3D simulation model based on finite element method for hepatic cryoablation planning. BMC Gastroenterology **25**(1), 227 (Apr 2025). https://doi.org/10.1186/s12876-025-03817-y,
14. Meister, F., Audigier, C., Passerini, T., Lluch, E., Mihalef, V., Maier, A., Mansi, T.: Fast Automatic Liver Tumor Radiofrequency Ablation Planning via Learned Physics Model. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2022, vol. 13437, pp. 167–176. Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-16449-1_17,
15. Nelder, J.A., Mead, R.: A Simplex Method for Function Minimization. The Computer Journal **7**(4), 308–313 (Jan 1965). https://doi.org/10.1093/comjnl/7.4.308,
16. Rieder, C., Schwenke, M., Pätz, T., Georgii, J., Ballhausen, H., Schwen, L.O., Haase, S., Preusser, T.: Evaluation of a numerical simulation for cryoablation – comparison with bench data, clinical kidney and lung cases. International Journal of Hyperthermia **37**(1), 1268–1278 (Jan 2020). https://doi.org/10.1080/02656736.2020.1845402, , _eprint: https://doi.org/10.1080/02656736.2020.1845402
17. Um, K., Brand, R., Yun, Fei, Holl, P., Thuerey, N.: Solver-in-the-Loop: Learning from Differentiable Physics to Interact with Iterative PDE-Solvers (Jan 2021). https://doi.org/10.48550/arXiv.2007.00016, , arXiv:2007.00016 [physics]
18. Xi, J., Lamata, P., Niederer, S., Land, S., Shi, W., Zhuang, X., Ourselin, S., Duckett, S.G., Shetty, A.K., Rinaldi, C.A., Rueckert, D., Razavi, R., Smith, N.P.: The estimation of patient-specific cardiac diastolic functions from clinical measurements. Medical Image Analysis **17**(2), 133–146 (Feb 2013). https://doi.org/10.1016/j.media.2012.08.001,