

Design of HIFU Treatment Plans using an Evolutionary Strategy

Marta Cudova
Centre of Excellence IT4Innovations,
Faculty of Information Technology,
Brno University of Technology
Brno, Czech Republic
icudova@fit.vutbr.cz

Bradley E. Treeby
Department of Medical Physics and
Biomedical Engineering,
University College London
London, United Kingdom
b.treeby@ucl.ac.uk

Jiri Jaros
Centre of Excellence IT4Innovations,
Faculty of Information Technology,
Brno University of Technology
Brno, Czech Republic
jarosjir@fit.vutbr.cz

ABSTRACT

High Intensity Focused Ultrasound (HIFU) is an emerging technique for non-invasive cancer treatment where malignant tissue is destroyed by thermal ablation. Since one ablation only allows a small region of tissue to be destroyed, a series of ablations has to be conducted to treat larger volumes. To maximize the treatment outcome and prevent injuries such as skin burns, complex preoperative treatment planning is carried out to determine the focal position and sonication time for each ablation. Here, we present an evolutionary strategy to design HIFU treatment plans using a map of patient specific material properties and a realistic thermal model. The proposed strategy allows high-quality treatment plans to be designed, with the average volume of mistreated and under-treated tissue not exceeding 0.1 %.

CCS CONCEPTS

• **Computing methodologies** → *Modeling methodologies*; • **Applied computing** → *Consumer health*;

KEYWORDS

Evolutionary strategy, HIFU, treatment planning, k-Wave.

ACM Reference Format:

Marta Cudova, Bradley E. Treeby, and Jiri Jaros. 2018. Design of HIFU Treatment Plans using an Evolutionary Strategy. In *GECCO '18 Companion: Genetic and Evolutionary Computation Conference Companion, July 15–19, 2018, Kyoto, Japan*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3205651.3208268>

1 INTRODUCTION

In recent years, High Intensity Focused Ultrasound (HIFU) has been applied to treat a variety of solid malignant tumors in a well-defined volume, including the pancreas, liver, prostate, breast, uterine fibroids, and soft-tissue sarcomas. In comparison to conventional tumor/cancer treatment modalities, such as open surgery, radio- and chemo-therapy, HIFU is non-invasive, non-ionising, and has fewer complications after treatment. Over 100,000 cases have been treated throughout the world with great success [21].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '18 Companion, July 15–19, 2018, Kyoto, Japan

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5764-7/18/07...\$15.00

<https://doi.org/10.1145/3205651.3208268>

The fundamental principles of HIFU ablation are coagulative thermal necrosis due to the absorption of ultrasound energy in the focal region, in addition to induced cavitation damage. When treatments are performed under MR-guidance, the thermal mechanism for tissue ablation is generally used, since this is better understood and its effect is easier to control. The basic principle of thermal HIFU treatment is to deliver sufficient ultrasound energy to raise the temperature by several tens of degrees so that the tissue is destroyed via coagulative necrosis. The HIFU beam focusing results in cytotoxic levels of temperature only at a specific location within a small volume (e.g., about 1 mm in diameter and about 10 mm in length), which minimizes the potential for thermal damage to tissue outside the focal region. The boundary between disrupted cells and normal tissue is typically less than 50 μm in width [17]. Large tumors can be destroyed by producing a contiguous lesion lattice encompassing the tumor and appropriate margins of surrounding tissue. However, complications may develop if vital blood vessels adjacent to the tumors are severely damaged. Moreover, blood perfusion may carry away a significant amount of energy and deteriorate the treatment outcome [9].

Despite the advantages of HIFU, the technique still suffers from long treatment times, and does not have the delivery precision of other established therapies such as radiotherapy. With recent advances in numerical methods and IT, detailed simulations accurately capturing the relevant physical behavior of focused ultrasound waves and temperature distribution in heterogeneous tissue are now possible [16]. This provides a very powerful tool that can be leveraged for a range of tasks, including patient selection (determining whether a patient is a good candidate for a particular procedure based on their individual anatomy) and treatment verification (determining the cause of adverse events or treatment failures). However, model-based treatment planning (determining the best transducer position and sonication parameters to deliver the ultrasound energy to the planning target volume) is still currently performed in a relatively rudimentary way based on heuristics rather than physical models of the therapy.

This paper presents a preliminary investigation into solving the treatment planning problem via an optimisation-based approach using evolutionary algorithms. The inputs to the optimisation are the planning target volume and organs at risk, and the evolutionary strategy produces a suitable treatment plan. Although the idea is quite simple, it has the potential to improve targeting ability and reduce treatment times, which will ultimately improve clinical outcomes for patients undergoing HIFU.

2 PROPOSED ALGORITHM

This section first describes the optimization algorithm based on the Matlab implementation of the Covariance Matrix Adaptation (CMA) Evolutionary Strategy (ES), version 3.61, developed by Nikolaus Hansen [4]. Then, the solution encoding is outlined. Finally, the fitness function based on the tissue-realistic heat distribution developed as part of the k-Wave toolbox is introduced [16, 18].

2.1 Evolutionary Strategy

This section provides a brief summary of the CMA-ES [4, 5] algorithm used for the design of HIFU treatment plans. The CMA [6] describes the pairwise dependencies between variables/genes on the top of the classic ES. The CMA-ES is a very popular stochastic method for real-parameter (continuous domain) optimization of nonlinear, nonconvex objective functions.

In the CMA-ES, a population of λ new search points (individuals, offspring) is generated by sampling a multivariate normal distribution $\mathcal{N}(\mathbf{m}, \mathbf{C})$ determined by its mean $\mathbf{m} \in \mathbb{R}^N$ and its symmetric and positive defined covariance matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$, which determines the shape of the distribution ellipsoid. The length of the step is controlled by the so-called step-size parameter $\sigma \in \mathbb{R}^N$:

$$x_i \sim \mathbf{m} + \sigma N_i(0, \mathbf{C}) \text{ for } i = 1, \dots, \lambda. \quad (1)$$

The newly generated individuals are first ranked according to their fitness and then the best μ individuals are selected. The elitism is not used. Next, the mean value, step size and the covariance matrix are updated. The mean value \mathbf{m} is updated by weighted intermediate recombination where the weight of every selected individual is proportional to its rank. The CMA-ES utilizes an evolution path to control the step size σ . Conceptually, the evolution path is the search path the strategy takes over a number of generation steps. It can be expressed as a sum of consecutive steps of the mean \mathbf{m} . The adaptation of the covariance matrix follows a natural gradient approximation of the expected fitness. The adaptation procedure first learns all pairwise dependencies between all variables. Then, it conducts a principle component (eigenvectors) analysis (PCA) of steps sequentially in time and space. Finally, a new rotated problem representation is determined using the Mahalanobis metric [2].

The main benefit of the CMA-ES is a very small population and fast convergence for real-valued problems compared to Genetic Algorithms (GA) [1] or Estimation of Distribution Algorithms (EDA) [11]. The step-size control facilitates fast (log-linear) convergence and possibly linear scaling with the dimension. The covariance matrix adaptation increases the likelihood of previously successful steps and can improve performance by orders of magnitude [5].

2.2 Solution Encoding

The ablation of large target areas using HIFU requires multiple sonications to effectively cover this area. The candidate solution I describes the trajectory the HIFU transducer follows in the tissue during the treatment. The treatment is not continuous but proceeds at precisely defined points in the tissue where the HIFU focus is placed. The number of sonications is limited to N , usually low tens. The amount of energy delivered during a single sonication is given by the length of the sonication t_{on} and the length of the subsequent cooling interval t_{off} . One sonication can thus be defined as a 4-tuple




Figure 1: Visualization of the focal region generated by a bowl HIFU transducer and its variation with distance.

S_i composed of two spatial coordinates of the beam focus (only 2D problems are considered), and sonication and cooling intervals t_{on} and t_{off} , respectively:

$$I = (S_1, S_2, \dots, S_N), \text{ where } S_i = (x_i, y_i, t_{\text{on}, i}, t_{\text{off}, i}) \quad (2)$$

Thus, the treatment planning problem is defined here as finding the optimum position and sonication times for the specified number of ablations to destroy the planning target volume while sparing the organs at risk.

2.3 Fitness Function

The evaluation of the quality of candidate treatment plans is composed of several stages. First, the heat deposition for every sonication is determined using the predicted shape and position of the ultrasound focus, and the sonication length. Second, a numerical thermal model is executed to calculate the temperature distribution in the domain during the treatment. The heat model is executed for the whole sequence of sonications in the treatment plan. Third, the treatment domain is examined and the areas exceeding the thermal dose threshold are identified. Finally, the correctly treated and mistreated regions are marked and the quality of the plan assessed.

2.3.1 Calculation of the Heat Deposition. The first step of the treatment plan evaluation is the calculation of heat deposition. Physically, the shape and position of the focus are determined by the nonlinear propagation of the ultrasound wave through heterogeneous tissue. The amount of energy deposited in the focus is proportional to the medium absorption. For accurate prediction of the heat deposition precise full-wave models [18], simplified forward wave propagation models [8], or very simple ray-tracing models [12] can be used. Unfortunately, the execution time of all these models (order of minutes) is prohibitive for applications in evolutionary optimization techniques.

Therefore, several simplifications have to be made. The first assumption is the centre of the focus can be placed at coordinates given by the sonication $S_i = [x_i, y_i]$. This can be achieved by precise electronic steering of the ultrasound beam by the HIFU transducer [19]. The second assumption is that the focus has an elliptical shape with the axes length proportional to the distance from the HIFU transducer, see Fig. 1. Since the distribution of the energy in the focus is known to follow the Gaussian distribution [3, 10, 20], we can define the ellipse axes by the distribution variance and set the magnitude of the volume rate of heat deposition according to




Figure 2: The temperature course in the focus centre during a sonication with $t_{on} = 20$ s and $t_{off} = 100$ s.

the power of the transducer (i.e., an acoustic model is not used to calculate the heat source term). For the sake of simplicity, we always assume the transducer axes to be aligned with the domain axes, and only consider problems in 2D. The total energy deposited in the tissue is then proportional to the time t_{on} .

3.2.2 Thermal Model Execution. The second step is the execution of the numerical thermal model to calculate the heat diffusion in the tissue. For modelling of heat diffusion the so-called Pennes’ bioheat equation [14] is applied, which has a corresponding ultrasound energy absorption term as a source term, see Eq. (3). For the energy balance in the heat transfer, it is of particular importance to take blood perfusion into account, which cools down the tissue and can prevent the thermal destruction of malignant cells close to blood vessels. In particular, with increasing heating duration, e.g., for volumetric ablation of large tissue areas, this effect will increase. The Pennes’ equation reads

$$\rho C \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) + W_b C_b (T - T_a) + Q \quad (3)$$

where C and C_b are specific heat of tissue and blood [J/kg°C], W_b [kg/m³s] is a blood perfusion related parameter, T_a is the arterial temperature (assumed to be 37 °C), and Q [W/m³] is the power deposited in the tissue by the ultrasound transducer.

Figure 2 shows the course of the temperature at the centre of the focus for a sonication period of 20 s and cooling period 100 s. After the transducer is turned on, the temperature starts to grow. Once the transducer is turned off (cooling period), the temperature slowly declines back to the ambient values due to the energy diffusion and blood perfusion.

Figure 3 shows the spatial temperature distribution along the main focus axis at the end of the sonication, and then every 20 s during the cooling period. Without the proximity of a large blood vessel, the spatial temperature distribution follows the Gaussian distribution. The temperature magnitude declines from a peak around 72 °C at the end of the sonication down to about 46 °C at the end




Figure 3: Heat distribution along the main focus axis during a sonication with $t_{on} = 20$ s and $t_{off} = 100$ s at time $t = 20$ s, 40 s, 60 s, 80 s, 100 s.

the cooling interval. On the other hand, the area with temperature exceeding 37 °C is slowly growing.

Thermal damage is usually computed using the Sapareto-Dewey iso-effect thermal dose relationship [15] which is expressed in seconds and represents the equivalent time which would produce the same biological effects at a temperature of 43 °C. This metric is called cumulative equivalent minutes at 43 °C (CEM₄₃). CEM₄₃ is calculated for every point in the tissue, see Eq. (4) and summed up over all sonications S_i .

$$CEM_{43} = \int_0^{t_{on}+t_{off}} R^{(43-T)} dt, \text{ where} \quad (4)$$

$$R = \begin{cases} 0 & \text{for } T \leq 39^\circ\text{C} \\ 0.25 & \text{for } 39^\circ\text{C} < T \leq 43^\circ\text{C}. \\ 0.5 & \text{for } T > 43^\circ\text{C} \end{cases}$$

Thermal doses of 240 minutes at 43 °C irreversibly damage and coagulate critical cellular protein, tissue structural components and the vasculature leading to immediate tissue destruction, however, the threshold varies with tissue type [21]. The area with the dosage exceeding 240 CEM₄₃ is depicted in Fig. 3 by a yellow bar.

The thermal model is implemented in Matlab using the k-Wave toolbox [18] and supports precise tissue parameter settings derived from patient-specific models of the tissue anatomy. The computation is based on a k-space pseudospectral scheme in which spatial gradients are calculated using the Fourier collocation spectral method, and temporal gradients are calculated using a k-space corrected finite difference scheme. For a homogeneous medium, the formulation is exact and unconditionally stable. For a heterogeneous medium, the time scheme allows larger time-steps to be taken for the same level of accuracy compared to conventional pseudospectral time-domain methods. The simulation domain is thus discretized into a grid with spatial and temporal resolution set according to the convergence testing.

2.3.3 Evaluation of the Treated Area. The output of the thermal model is a spatial map of CEM_{43} cumulated over the whole treatment I . This map is thresholded by a value of 240 to produce a binary mask of destroyed tissue.

The evaluation of the quality of the HIFU treatment is based on the assumption that all tissue in the target area is destroyed while all tissue in the prohibited area (organs at risk) is left unharmed. In order to give the optimization algorithm some freedom, don't care areas can be specified as well.

The fitness function for a 2 dimensional case can be written as

$$f = \int_0^X \int_0^Y ((D * \bar{C}) + (P * C)) dx dy$$

$$C = \begin{cases} 0 & \text{for } CEM_{43} \leq 240 \\ 1 & \text{for } CEM_{43} > 240 \end{cases} \quad (5)$$

$$D \in \mathbb{R}^+, P \in \mathbb{R}^+,$$

where X, Y are the sizes of the domain along the x and y dimension, respectively, C is the binary mask representing the actually treated area, \bar{C} is the complementary mask representing the non-treated area, D is the target map specifying the area to be treated and P is the prohibited area. Since D and P are defined as functions over a 2D space, the user can specify the level of urgency a given point in the space shall be treated or protected with. This enables two substantial concepts to be incorporated in the fitness function: (1) a fine-grained tuning of the shape of the target and prohibited areas, (2) introduction of gradient into the fitness function space which guides the evolutionary process away from prohibited areas towards the desired areas. The goal is then to minimize the fitness function.

3 EXPERIMENTAL RESULTS

The experimental work presented in this paper has three main goals: (1) confirm the hypothesis that ES is able to design suitable HIFU treatment plans, (2) find suitable parameters of the ES, and (3) evaluate the computational complexity of the ES.

3.1 HIFU Treatment Setup

During the evaluation of the proposed algorithm we strove to work under as realistic conditions as possible. Thus, the material properties were derived from a realistic patient dataset, the transducer parameters from a realistic HIFU transducer and the simulation parameters set according to a numerical study.

3.1.1 Material Properties. To demonstrate the ability of the proposed algorithm to design suitable treatment plans, a representative map of the material properties in the human body is required. Here, the open-source AustinWoman voxel model is used [13]. This is a segmentation of a digital image dataset acquired as part of the Visible Human Project run by the U.S. National Library of Medicine.

The original dataset consists of digital cryosection images with 0.33 mm pixel spacing taken at 0.33 mm intervals through the axial direction of a 59-year-old female cadaver. Radiological images (CT and MRI) are also available. The AustinWoman segmentation divides the digital image dataset into 58 material labels that cover the major tissue types present in the body. The tissue density, specific heat and thermal conductivity are then assigned using the average

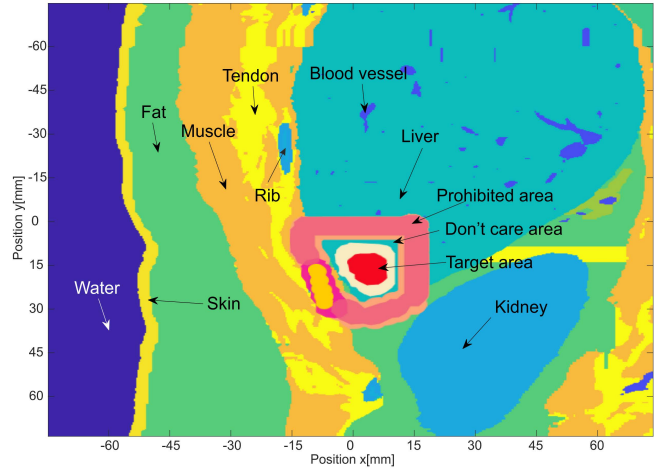


Figure 4: Segmented AustinWoman data with marked target and prohibited areas.

properties from the IT'IS tissue property database V3.0 [7]. For the sake of simplicity, capillary blood perfusion is only considered during this evaluation.

As a case study for HIFU treatment planning, one abdominal target within the right lobe of the liver is used. Figure 4 shows the segmented part of the body used for the treatment planning. The target area is marked close to the bottom of the liver. Two levels of D are chosen, a higher one ($D_{x,y} = 2$) in the middle of the target area, and a lower one ($D_{x,y} = 1$) close to the boundaries. The prohibited area P is marked by three different levels of importance. The highest one ($P_{x,y} = 5$) covers the rib and the tendon. The middle one ($P_{x,y} = 2$) covers the fat layers and the areas further from the treated areas. The lowest level ($P_{x,y} = 1$) is used for areas neighbouring the treated areas. To make the interface between treated and prohibited areas smooth, a thin don't care area is used ($D_{x,y} = 0$ and $P_{x,y} = 0$).

3.1.2 Transducer Properties. The size of the heat source was based on a single element transducer using the nominal properties of the HAIFU JC-200¹. This has a radius of curvature of 145 mm, aperture diameter of 200 mm, and frequency of 0.95 MHz, giving a focal beam ellipse size of approximately 19×2.5 mm. This size was used to define the full-width at half maximum of the ellipsoidal heat source term. The spatial peak of the volume rate of heat deposition was set to 100 W/cm², which approximately matches the values used for clinical treatments.

3.1.3 Numerical Model Parameters. The parameters of the numerical thermal model were set according to numerical convergence testing as follows:

- Discretized simulation domain size 495 × 495 grid points, periodic boundary condition.
- Spatial resolution 0.2 mm. The original AustinWoman data was upsampled using linear interpolation.
- Temporal resolution 0.1 s.

¹<http://http://www.haifummedical.com>

- The total length of the simulation $\sum_{i=0}^N (t_{on,i} + t_{off,i})$.
- Allowed positions of the ultrasound focus centre limited to the bounding box at grid positions $[270, 230] \times [345, 295]$.
- Maximum sonication and cooling periods $t_{on} = [0, 20 \text{ s}]$, $t_{off} = [0, 100 \text{ s}]$.
- Number of sonications considered $N \in \{4, 5, 6, 8, 10\}$.

The numerical model as well as the ES were executed in Matlab 2015b on the Anselm supercomputer². The execution was parallelized using the Matlab Parallel Computing Toolbox over 16 processor cores.

3.2 Optimization Parameters Setting

The CMA-ES does not require a tedious parameter tuning for its application. The default parameter values were used as follows:

- The default population size λ is comparatively small to allow fast convergence and equals to $\lambda = 4 + \lfloor (3 * \ln(4 * N)) \rfloor$, where N is the number of sonications, and $4 * N$ the number of variables in the individual. Numerically expressed, $\lambda = 12, 12, 13, 14, 15$ for $N = 4, 5, 6, 8, 10$, respectively.
- The size of parent population $\mu = \lfloor (\lambda/2) \rfloor$.
- The step size for every variable $\sigma_i = (U_i - L_i)/3$, where U_i and L_i are the upper and lower bounds for a given variable.

3.3 Evaluation of the Evolution Process

The ability of the CMA-ES to find a suitable treatment plan for a given patient is evaluated in this section. To allow some statistical evaluation, 15 independent runs were executed with the maximum execution time per run of 48 hours. We examine three basic metrics: (1) the quality of the best treatment plan, (2) the number of fitness evaluations to converge, and (3) the development of the best individual in time.

3.3.1 Quality of the Treatment Plans. The most important criteria is the ability of the CMA-ES to find a good treatment plan. The perfect solution, with the fitness value equalling to 0, covers the whole treated area while leaves the prohibited area unharmed.

Figure 5 shows the histogram of the fitness values collected over the best solutions generated in 15 independent runs. Here, the influence of the population size as well as the number of sonications on the quality of the treatment plan is shown.

The first observation indicates that the higher the number of sonications the CMA-ES is allowed to use, the higher the average quality of the treatment plan is achieved. Naturally, the more (potentially shorter/smaller) sonications, the more precise coverage control. On the other hand, the total treatment becomes longer, which may make the treatment more difficult for the patient, and for the computer to evaluate the plan (more runs of the thermal model). Yet, it is possible to design an optimal treatment plan even with 4 sonications, provided a suitable population size is chosen. Nevertheless, the success rate of finding an optimal treatment plan is quite low, even for the biggest population of 40 individuals.

If the stress on the minimal number of sonications is not so strong, the CMA-ES is able to produce optimal treatment plans in most of its runs. The best trade-off between the treatment plan quality and the number of sonications seems to be 8 sonications,




Figure 5: Histogram of the best fitness values from 15 independent runs as a function of the population size and the number of sonications.

and the population size of 20 or 40 individuals. Finally, the quality appears to deteriorate for 10 sonications and large population sizes. This is caused by an extended time to evaluate individuals. Together with a large population, the computational resources (48 hours of wall clock time) are depleted quite quickly without reaching the convergence.

Since the fitness value is an aggregation of two metrics, see Eq. (5), we also provide separate metrics for the treatment plan quality. Tables 1 and 2 present the median and average percentage of the non-treated and mistreated areas, respectively. The non-treated area marks places where the treatment was not successful (the energy deposition was too low) while the mistreated area marks

Table 1: Median percentage of non-treated/mistreated area from 15 independent runs as a function of the population size λ and the number of sonications.

λ	Number of Sonications				
	4	5	6	8	10
def	0.03/0.10	0.02/0.06	0.02/0.04	0.00/0.00	0.00/0.00
20	0.00/0.11	0.00/0.09	0.00/0.01	0.00/0.00	0.00/0.00
40	0.00/0.09	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00

Table 2: Average percentage of non-treated/mistreated area from 15 independent runs as a function of the population size λ and the number of sonications.

λ	Number of Sonications				
	4	5	6	8	10
def	0.10/0.10	0.03/0.05	0.05/0.05	0.00/0.02	0.00/0.01
20	0.04/0.10	0.01/0.08	0.00/0.04	0.00/0.01	0.00/0.01
40	0.04/0.07	0.00/0.03	0.00/0.03	0.00/0.01	0.02/0.00

²IT4Innovations National Supercomputing Centre, Czech Republic

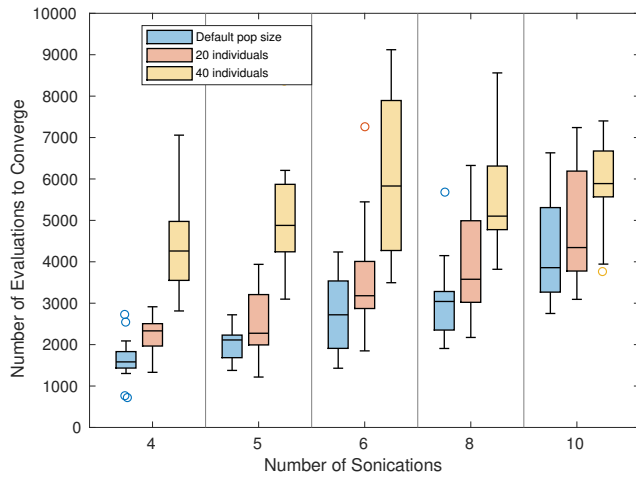


Figure 6: Number of evaluations to reach convergence for various numbers of sonications and population size.

places where the prohibited area was damaged. The key message of both tables is that even if the treatment plan is not perfect, the non-treated and mistreated areas comprise less than 0.1 % of the marked regions. We consider this as an exceptional results, which is fully human-competitive.

Second, since the medians of both metrics equal to zero for 8 and 10 sonications regardless of the population size, we can conclude more than 50 % of runs produced perfect treatment plans. In many cases, it also appears that it is easier to cover the treated area than to protect the prohibited area. This is given by the energy leaking from the focus to the protected area (heat diffusion).

3.3.2 Convergence Speed. The second important factor determining the quality of the proposed evolutionary strategy is the convergence speed. This metric can be measured in terms of the number of fitness function evaluations (see Fig. 6), or the wall clock time (see Table 3) needed to find either an optimal solution or to converge to a suboptimal solution.

Figure 6 presents a histogram of the number of fitness function evaluations as a function of the number of sonications and the population size. First, the population size clearly has a significant impact on the number of evaluations. This is particularly noticeable for the population size of 40 individuals, where the CMA-ES needs more than twice as many evaluations to converge as with the

Table 3: Median of the execution time of the Matlab implementation running on a 16-core node of the Anselm super-computer measured over 15 runs as a function of the population size λ and the number of sonications.

λ	Number of Sonications				
	4	5	6	8	10
default	11.11h	16.02h	20.09h	23.78h	31.62h
20	15.20h	19.36h	28.07h	28.26h	32.22h
40	25.80h	31.78h	43.11h	33.87h	46.96h




Figure 7: Typical progress of the evolution for 8 sonications with 20 individuals.

population of 20 individuals. Unfortunately, the gain in the solution quality cannot justify the growth in the evolution time. It appears higher diversity cannot help with finding a perfect solution much.

On the other hand, increasing the number of sonications has only a small impact on the number of evaluations. This can be attributed to the quality of the CMA-ES which is said to have linear scaling with the problem dimension [4].

3.3.3 Evolution Progress. Finally, we investigated the progress of the CMA-ES by observing the evolution of the best individual in the population. Figure 7 shows the statistics of the development of the best individual over 15 independent runs with 8 sonications and 20 individuals on a logarithmic scale. It can be seen that a typical run converges quite rapidly and 250 generations are usually enough to converge.

3.4 Visualization of Treatment Plans

Here, the visualization of two treatment plans is presented, a perfect one in Fig. 8, and a suboptimal in Fig. 9. Each row in those figures represents an individual sonication in the order it appears in the treatment plan. The first column shows the actual temperature distribution in the domain after the sonication (including the cooling period). We can see that part of the energy from previous sonications has not disappeared yet and is being used by the following sonications to raise the temperature. The second column shows the lesion map, more precisely the areas where CEM_{43} has exceeded a value of 240 mins. The third column shows the areas which have not been treated properly yet ($CEM_{43} < 240$). In Fig. 8, the whole target area is covered, eventually. On the other hand, Fig. 9 shows that even after all four sonications, there is a small area of tissue that has not received the target dose of 240 CEM. The last column shows the mistreated area. In the first case, there is no mistreatment, however, in the second case, a small part of the protected tissue was damaged. It is left for the clinician decision to evaluate the significance these under-treatments and mistreatments.

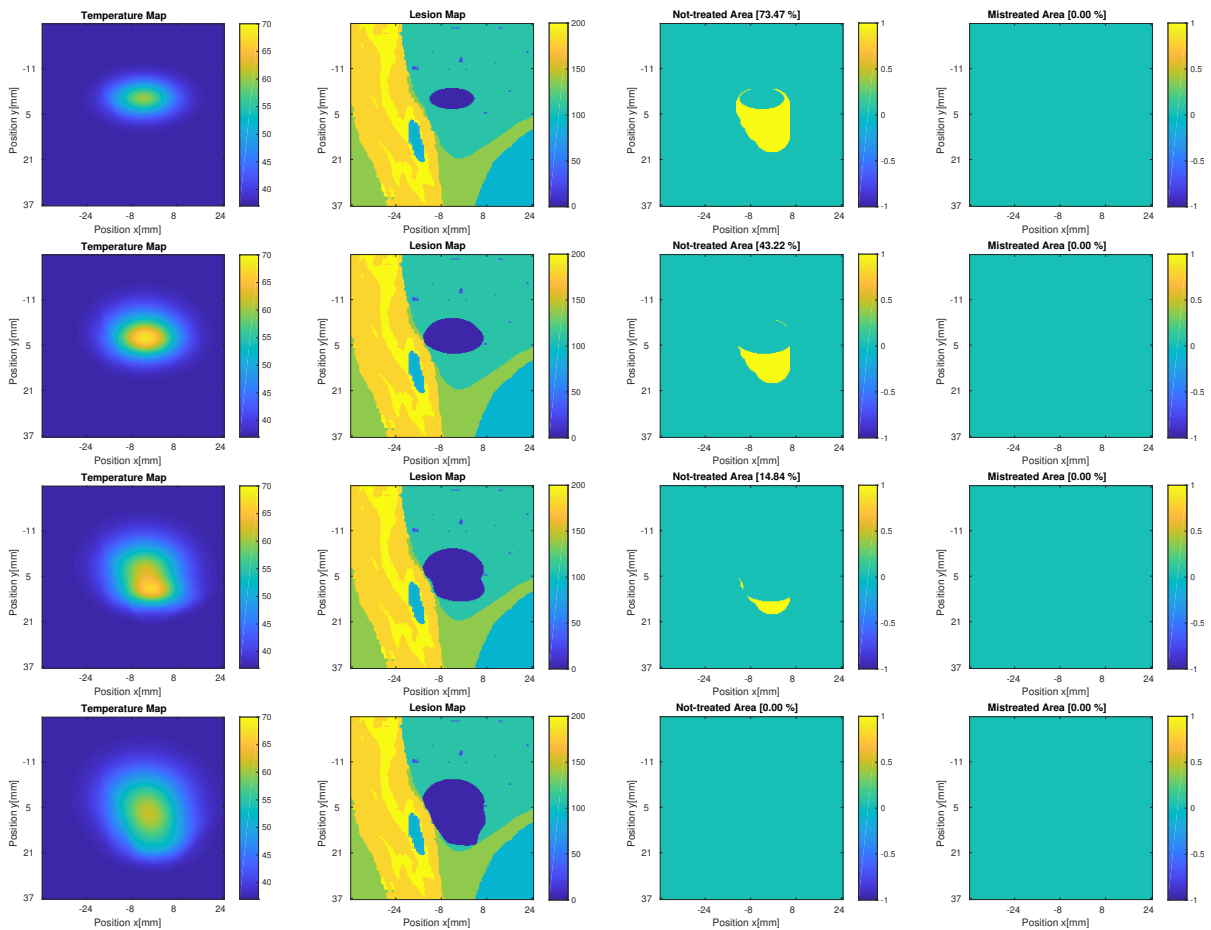


Figure 8: Visualization of a perfect treatment plan composed of 4 sonications.

4 CONCLUSION

High Intensity Focused Ultrasound (HIFU) is an emerging technique which has a great potential in non-invasive surgery and cancer treatment. However, to precisely target the ultrasound beam and destroy only the tumor and leave the healthy tissue untouched, complex preoperative treatment planning is required. However, manual and semi-automated planning are very tedious and computationally intensive jobs. Thus, new approaches for fully automated design of treatment plans are of a great interest.

This paper serves as a proof of concept showing that it is possible to use an evolutionary strategy to design close to optimal treatment plans with a negligible mistreated and undertreated areas on the order of 0.1%. The great benefit of this algorithm is that only a map of the patient anatomy and the corresponding target and prohibited areas are required. The rest is handled by the evolution.

On the other hand, the presented approach has its limitations. First, it usually takes between 36 to 48 hours to design a good treatment plan, even when an acoustic model is not used, and the thermal simulations are conducted in 2D. This, however, is going to be addressed as the next step in our research, which is the reimplementation of the whole algorithm in high performance languages

such as parallel C/C++ with the aim to reduce the computational time by a factor of 5, at least. Second, we had to limit ourselves to only work in two dimensions due to computational complexity. Third, we would like to develop a model coupling between the ultrasound and thermal modes to account for acoustic heterogeneities in the designed treatment plans.

5 ACKNOWLEDGEMENT

This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic from the National Programme of Sustainability (NPU II) project IT4Innovations excellence in science - LQ1602, and from the Large Infrastructures for Research, Experimental Development and Innovations project IT4Innovations National Supercomputing Center - LM2015070. This project has received funding from the European Union's Horizon 2020 research and innovation programme H2020 ICT 2016-2017 under grant agreement No 732411 and is an initiative of the Photonics Public Private Partnership. This project was also supported by the Engineering and Physical Sciences Research Council, UK, grant numbers EP/L020262/1 and EP/M011119/1.

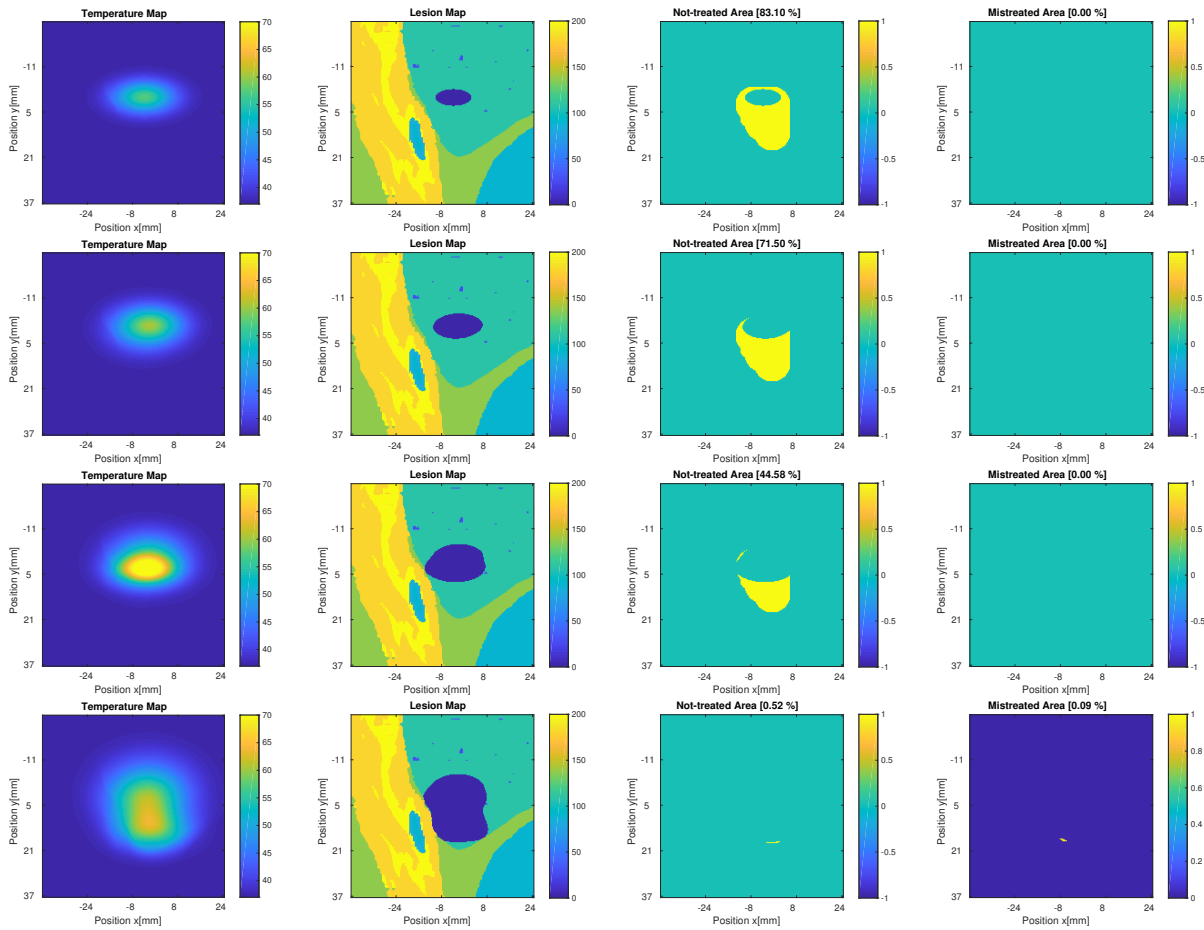


Figure 9: Visualization of a suboptimal treatment plan composed of 4 sonications with a fitness value of 30.

REFERENCES

[1] E. Cantú-Paz. 2000. *Efficient and Accurate Parallel Genetic Algorithms*. Kluwer Academic Publishers. 187 pages.

[2] S. Chen, B. Ma, and K. Zhang. 2009. On the similarity metric and the distance metric. *Theoretical Computer Science* 410, 24-25 (2009), 2365–2376.

[3] C.R. Dillon, U. Vyas, A. Payne, et al. 2012. An analytical solution for improved HIFU SAR estimation. *Physics in Medicine and Biology* 57, 14 (2012), 4527–4544.

[4] N. Hansen. 2006. The CMA Evolution Strategy: A Comparing Review. In *Towards a New Evolutionary Computation*. Vol. 192. Springer-Verlag, Heidelberg, 75–102.

[5] N. Hansen. 2016. *The CMA Evolution Strategy: A Tutorial*. Technical Report. arXiv:1604.00772

[6] N. Hansen and A. Ostermeier. 1996. Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In *Proceedings of IEEE International Conference on Evolutionary Computation*. IEEE, 312–317.

[7] P. Hasgall, F. Di Gennaro, C. Baumgartner, et al. 2015. IT’IS Database for thermal and electromagnetic parameters of biological tissues, version 3.0. (2015). <https://doi.org/10.13099/VIP21000-03-0>

[8] J. Huijssen and M.D. Verweij. 2010. An iterative method for the computation of nonlinear, wide-angle, pulsed acoustic fields of medical diagnostic transducers. *The Journal of the Acoustical Society of America* 127, 1 (jan 2010), 33–44.

[9] M. Ichihara, K. Sasaki, S.I. Umemura, et al. 2007. Blood flow occlusion via ultrasound image-guided high-intensity focused ultrasound and its effect on tissue perfusion. *Ultrasound in Medicine and Biology* 33, 3 (2007), 452–459.

[10] J.W. Jenne, T. Preusser, and M. Günther. 2012. High-intensity focused ultrasound: Principles, therapy guidance, simulations and applications. *Zeitschrift für Medizinische Physik* 22, 4 (dec 2012), 311–322.

[11] J. Madera and B. Dorronsoro. 2006. Estimation of Distribution Algorithms. *Metaheuristic Procedures for Training Neural Networks* December (2006), 87–108.

[12] F. Marquet, M. Perndoiot, J.-F. Aubry, et al. 2009. Non-invasive transcranial ultrasound therapy based on a 3D CT scan: protocol validation and in vitro results. *Physics in Medicine and Biology* 54, 9 (2009), 2597–2613.

[13] J.W. Massey, C.S. Geyik, N. Techachainiran, C. Hsu, R.Q. Nguyen, T. Latson, M. Ball, E. Celik, and A.E. Yilmaz. 2012. AustinMan and AustinWoman: High fidelity, reproducible, and open-source electromagnetic voxel models. In *34th Meeting of the Bioelectromagnetics Society*.

[14] H. Pennes. 1948. Analysis of tissue and arterial blood temperatures in the resting human forearm. *Journal of applied physiology* 1, 2 (1948), 93–122.

[15] S.A. Sapareto and W.C. Dewey. 1984. Thermal dose determination in cancer therapy. *Int. Journal of Radiation Oncology, Biology, Physics* 10, 6 (1984), 787–800.

[16] V. Suomi, J. Jaros, B.E. Treeby, and R. Cleveland. 2017. Full Modelling of High-Intensity Focused Ultrasound and Thermal Heating in the Kidney using Realistic Patient Models. *IEEE Transactions on Biomedical Engineering* PP, 99 (2017), 1–12.

[17] G. Ter Haar, I. Rivens, L. Chen, and S. Riddler. 1991. High intensity focused ultrasound for the treatment of rat tumours. *Physics in Medicine and Biology* 36, 11 (1991), 1495–1501.

[18] B.E. Treeby, J. Jaros, A.P. Rendell, and B.T. Cox. 2012. Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method. *The Journal of the Acoustical Society of America* 131, 6 (2012), 4324–36.

[19] J. Vincenot, D. Melodelima, F. Chavrier, et al. 2013. Electronic Beam Steering Used with a Toroidal HIFU Transducer Substantially Increases the Coagulated Volume. *Ultrasound in Medicine and Biology* 39, 7 (2013), 1241–1254.

[20] G. Ye, P.R. Smith, and J.A. Noble. 2010. Model-Based Ultrasound Temperature Visualization During and Following Hifu Exposure. *Ultrasound in Medicine and Biology* 36, 2 (2010), 234–249.

[21] Y-F. Zhou, A. Syed Arbab, and R.X. Xu. 2011. High intensity focused ultrasound in clinical tumor ablation. *World journal of clinical oncology* 2, 1 (2011), 8–27.