

An Artificial Intelligence Dose Engine for Fast Carbon Ion Treatment Planning



A. Quarz^{1,2}, A. De Gregorio³, G. Franciosini^{4,5}, A. Schiavi^{4,5}, Z. Perko⁶, L. Volz¹, C. Hoog Antink², V. Patera^{4,5}, M. Durante^{1,7,8}, C. Graeff^{1,2,*}

¹ GSI Helmholtzzentrum für Schwerionenforschung, Biophysics Department, Darmstadt, Germany

² Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, Darmstadt, Germany

³ Dipartimento di Fisica, Sapienza Università di Roma, Roma, Italy

⁴ Dipartimento di Scienze di Base e Applicate per l'Ingegneria, Sapienza Università di Roma, Roma, Italy

⁵ INFN (Istituto Nazionale di Fisica Nucleare) section of Roma 1, Rome, Italy

⁶ Department of Radiation Science and Technology, Delft University of Technology, Delft, The Netherlands

⁷ Department of Condensed Matter Physics, Technische Universität Darmstadt, Darmstadt, Germany

⁸ Department of Physics "Ettore Pancini", University Federico II, Naples, Italy

ARTICLE INFO

Keywords:

Treatment planning
Dose engine
RBE-weighted dose
Deep-learning
Carbon therapy

ABSTRACT

Purpose: Monte Carlo (MC) simulations provide gold-standard accuracy for carbon ion therapy dose calculations but are computationally intensive, limiting their use in adaptive workflows. Analytical pencil beam algorithms offer speed but reduced accuracy in heterogeneous tissues. This study develops the first AI-based dose engine capable of predicting relative biological effectiveness-weighted doses. Absorbed dose, α , and β parameters for optimization are calculated at MC-level accuracy with a drastically reduced computational time.

Materials and Methods: We extended the transformer-based DoTA architecture to predict absorbed dose (C-DoTA-d), α (C-DoTA- α), and β (C-DoTA- β), introducing a cross-attention mechanism for α and β to combine dose and energy inputs. The training dataset consisted of approximately 70 000 pencil beams from 187 head-and-neck patients, with ground-truth values obtained using the GPU-accelerated MC toolkit FRED. Performance was evaluated on an independent test set using gamma pass rate (1%/1 mm), depth-dose, and isodose contour Dice coefficients. MC dropout-based uncertainty analysis was performed.

Results: Median gamma pass rates exceeded 98% for all predictions (99.76% for dose, 99.14% for α , 98.74% for β), with minima above 85% in the most heterogeneous anatomies. The Dice coefficient was 0.95 for 1% isodose contours, with slightly reduced agreement in high-gradient regions. Compared to MC FRED, inference was over $400\times$ faster (0.032 vs 14 seconds per pencil beam) while maintaining accuracy. Uncertainty analysis showed high stability, with mean standard deviations below 0.5% for all models.

Conclusions: This AI-based dose engine achieves MC-quality predictions of absorbed dose and relative biological effectiveness model parameters in ~ 30 ms per beamlet. Its speed and accuracy support online adaptive planning, paving the way for more effective carbon ion therapy workflows. Future work will expand to additional anatomical sites, beam geometries, and clinical beamlines.

Introduction

Particle therapy beams exhibit a steep maximum at the end of their range in the patient, known as the Bragg peak. Calculating this sharp dose gradient in heterogeneous tissue is particularly complex. For carbon ion therapy, the ions' variable relative biological effectiveness (RBE)—which also shows a maximum at the Bragg peak location—introduces a further

challenge. Accurate dose calculation in carbon ion therapy, therefore, not only requires a precise dose model, but it also relies on the correct calculation of biological factors that depend on the spectrum of primary and secondary particles along the beam path in the patient.^{1,2}

Monte Carlo (MC) particle transport simulations are generally considered the gold standard for dose calculation and radiobiological modeling.³ By considering the physics processes involved in the beam

* Corresponding author.

E-mail address: c.graeff@gsi.de (C. Graeff).

interaction with the patient, MC simulations accurately produce the lateral broadening and range mixing of the pencil beams (PBs). Recent advances in GPU-accelerated MC codes have allowed increasing introduction of MC dose engines in commercial TPSs for proton therapy.^{4–6} However, the increased complexity of particle transport for carbon ions compared to protons—requiring the simulation of multiple fragment species along the primary ion path and the associated computational expense still render MC carbon dose engines impractical for clinical routine. Computation time for even the fastest currently available carbon ion GPU MC codes is a significant bottleneck, particularly in scenarios requiring many or rapid plan optimizations, such as for robust 4D optimization⁷ and adaptive treatment workflows.^{8,9}

Due to this limitation, all carbon radiotherapy treatments are still planned using variations of analytical pencil beam algorithms (PBAs).^{10,11} However, these algorithms make inherent assumptions that drastically limit their accuracy in heterogeneous tissues.¹² Several studies have compared PBAs with MC methods, highlighting trade-offs between speed and precision in proton¹³ and carbon therapy.^{13,14} For high-precision treatment planning that fully exploits the accuracy available with carbon ion beams, MC-quality dose engines are essential.

For proton dose calculation, artificial intelligence (AI) has been successfully applied to achieve MC accuracy at massively reduced runtimes, sometimes even outperforming traditional PBA speed. Approaches such as DoTA,¹⁵ LSTM-based dose prediction,^{16,17} and DiscoGAN¹⁸ have demonstrated the feasibility of rapid dose calculations of single PBs as well as for the full dose distribution.¹⁹ Recent studies have also explored the use of AI to enhance the quality of PBA for protons.^{20,21}

The greater complexity of carbon ions due to nuclear fragmentation and their effect on the resulting RBE parameters present a challenge that has not been considered in any of the available AI dose engines yet. As any carbon therapy treatment planning system (TPS) needs to optimize the RBE-weighted dose, carbon ion dose engines must include the relevant biological parameters used in the RBE-weighted dose calculations, namely the alpha and beta parameters in the linear quadratic model of cell survival.¹¹

Previous works in carbon ion dose prediction^{22,23} have focused solely on physical dose, omitting the biologically crucial RBE component. To our knowledge, this study represents the first attempt to fill in this gap by developing an AI-driven dose engine capable of jointly predicting physical dose and α and β parameters for RBE-weighted optimization.

We extended the transformer-based DoTA¹⁵ model, originally developed for proton physical dose calculation, and trained 3 models separately to predict the input parameters for the local effect model (LEM I)²⁴: namely, the absorbed dose (C-DoTA- d), alpha (C-DoTA- α) and beta (C-DoTA- β). Our model operates on a per-PB basis, enabling both optimization and forward dose calculation.

In this proof-of-concept demonstration, we assess our model's feasibility on a cohort of head and neck cancer patients formerly treated within a carbon ion therapy pilot study at GSI (Darmstadt, Germany). As a baseline for comparison and to generate training data, we utilize the GPU MC toolkit FRED,^{4,25} which can score the alpha and beta parameters alongside the physical dose at fast per-pencil-beam calculation speeds of only a few seconds. The models deliver all 3 input parameters in only 32 ms, with a median gamma pass rate (GPR) of a gamma analysis with 1%/1 mm criteria and a 10% threshold > 99%.

Materials and methods

Patient data

Treatment planning data were taken from the GSI carbon ion pilot project for a cohort of head and neck cancer patients treated between 1997 and 2008 at GSI with carbon ions. The pilot project was approved by the ethical committee of the University of Heidelberg in 1997.

Anonymized treatment plans of all patients are stored for research purposes at GSI. Informed consent is waived by the ethical committee of the University of Heidelberg for anonymized plans used for research purposes.

For this study, 225 head-and-neck cancer patients were randomly selected from the database. 187 patients' CTs were used for training and validation of the models, and 38 for the final test. For each patient, the target contours and isocenters of the clinical plans were reused for PB data generation, but to augment the data, additional beam angles were added.

Data generation

As most patients were treated with lateral opposing fields from GSI's fixed horizontal beam line, also the immobilization masks were created for these angles. Therefore, for each patient, 5 positive and 5 negative couch angles were randomly selected from the range $+/-[85^\circ, 105^\circ]$. Treatment fields with a PB spacing of 10 mm were generated in TRiP98 for these angles, and a beam's-eye-view (BEV) CT was generated for each of the unique PB coordinates.

All BEV CTs were set to an isotropic resolution of 1 mm, with lateral slices of 48x48 pixels centered on the PB coordinate. BEV CT slices extended through the entire CT. Proximal slices containing only air ($HU < 900$) were removed, see Figure S1. BEV CTs that contained part of the shoulder or mask mounting structure were excluded.

Beam energy values in the dataset were restricted to the range 115 to 260 MeV/u to ensure a homogeneous representation of energies. For each BEV CT, two PB are selected, with an energy randomly sampled from the available discrete beam energies. If within this energy range, no Bragg Peak would be placed inside the target, this BEV CT was excluded. To achieve an even distribution, energies are excluded from sampling if the number of PB for this energy exceeds 99% of the maximum number of PB for any energy. If the Bragg Peak of the selected PB would be placed outside the patient, the energy was re-sampled. The final distribution of energies is shown in Figure S2.

The data generation resulted in 34 706 and 1769 BEV-CTs for training and validation, respectively (corresponding to a 95%/5% split), with 2 PB of different energy each. For the test dataset, from the remaining 38 patients 6963 BEV CTs with 13 926 PBs were generated. For each PB, ground truth (GT) data was computed with the GPU-based MC tool FRED.^{4,25}

FRED

A fast MC particle transport software was needed to compute the large amount of required GT data. We chose FRED,^{4,25} which is capable of GPU-accelerated dose calculations for protons and light ions. This work relies on the recently implemented carbon ion extension of FRED.^{4,25}

The simulation for carbon ions considers 3 main physics blocks: ionization energy loss, multiple scattering (both implemented analogously to the proton transport models of FRED), and a phenomenological fragmentation model. The latter was developed directly from published fragmentation data and evaluated against this data as well as FLUKA simulations. Full details on the physics processes in the simulation are provided in the overview paper by Simoni et al.⁴ In addition to the absorbed dose, FRED supports RBE-weighted dose calculation with multiple scoring options. In this work, we relied on table-based RBE calculation using tabulated alpha and beta values as a function of energy for the first 8 light ion species (up to Oxygen) generated with LEM I. Internally in FRED, the effect of mixed ion fields is resolved by dose averaging,^{11,26} in the same way as in TRiP98:

$$\alpha = \frac{\sum_{i=1}^n a_i D_i}{\sum_{i=1}^n D_i} \quad (2.1)$$

$$\beta = \left(\frac{\sum_{i=1}^n \sqrt{\beta_i} D_i}{\sum_{i=1}^n D_i} \right)^2 \quad (2.2)$$

Where α_i and β_i are the energy-specific alpha and beta values for a specific ion type, n is the total number of ion species considered. D_i is the dose contribution from each ion species. The mixed alpha and beta values were saved together with the absorbed dose. In the TRiP98 TPS, these values are used alongside the physical dose to calculate the RBE-weighted dose. Hence, they present the target values for the C-DoTA- α and C-DoTA- β models, respectively, while the absorbed dose output was the target for the C-DoTA-d model.

The simulation geometry consisted of the interpolated BEV CT, positioned so that its proximal entrance was located at the isocenter. The beam nozzle was simulated as a polymethyl methacrylate slab of 4 mm thickness, positioned 10 cm in front of the BEV CT in the beam direction. The beam was generated 20 cm in front of the isocentre. Each PB was simulated as a Gaussian beam with full width half maximum according to the CNAO beam characteristics at the respective PB energy. For all PBs, a momentum spread of 0.5% was assumed. Note that this does not reflect the characteristics of the CNAO beamline. A fully realistic simulation would introduce slight differences in the form of the beam's lateral spread (wider due to air drift) and range straggling (slightly broader). The simplified beam geometry was chosen to permit efficient training data generation. Each PB simulation used 10^7 primaries for high-quality GT data. Dose, alpha, and beta were scored on the same grid as the BEV CT. FRED provides the dose in Gy/primary.

Model architecture

Figure 1 represents the C-DoTA architectures. For physical dose prediction, the vanilla DoTA architecture proved optimal for carbon ions, utilizing normalized BEV CT scans and energy information as inputs (C-DoTA-d network). However, for α and β predictions, we extended the model by incorporating:

1. An additional convolutional encoder and transformer block processing both energy data and the physical dose output from the C-DoTA-d network (using a 1% dose threshold below which doses were zeroed out),
2. A cross-attention block to merge the BEV CT and dose information streams, and
3. A final decoder block for prediction output. The resulting models are called C-DoTA- α and C-DoTA- β , respectively.

Other model design choices in the original DoTA, such as the choice of convolutional and transformer block setups, hyperparameters (including learning rate schedules), and the optimizer configuration, were found to be optimal for this study.

Data preprocessing

Each PB simulated in FRED was cropped symmetrically to $24 \times 24 \times 240 \text{ mm}^3$. To prevent training on the noise at the beam penumbra, dose values below a given threshold were masked and zeroed out. For training of C-DoTA-d, this threshold was defined as 1% of the sample maximum in the FRED-computed absorbed dose. For consistency between generated training data and the intended application of the model, for C-DoTA- α and C-DoTA- β , a threshold of 1% of the sample maximum from the C-DoTA-d prediction was used instead (as during inference only the C-DoTA-d predicted is available). For distinction, the 2 masks are called "MC mask" and "AI mask" for the remainder of the document. This workflow is illustrated in Figure S3.

Unlike the vanilla DoTA model, the C-DoTA models normalized only the inputs, such as CT data and energy, and the absorbed dose predicted by C-DoTA-d, range to [0,1]. Normalization of the labels caused blurry output, especially near the Bragg peak. Instead, we used the following FRED labels in our loss function: absolute absorbed dose values, alpha scaled by 10 and beta values by 10^3 . This resulted in value ranges for the absorbed dose of [0-2.6 Gy], for alpha [0-15.8 Gy⁻¹], and for beta [0-26.5 Gy⁻²]. This scaling increased the values of the loss function and enabled more effective updates of the trainable weights. Moreover, interpretable thresholds could be employed in the custom loss, see the following section.

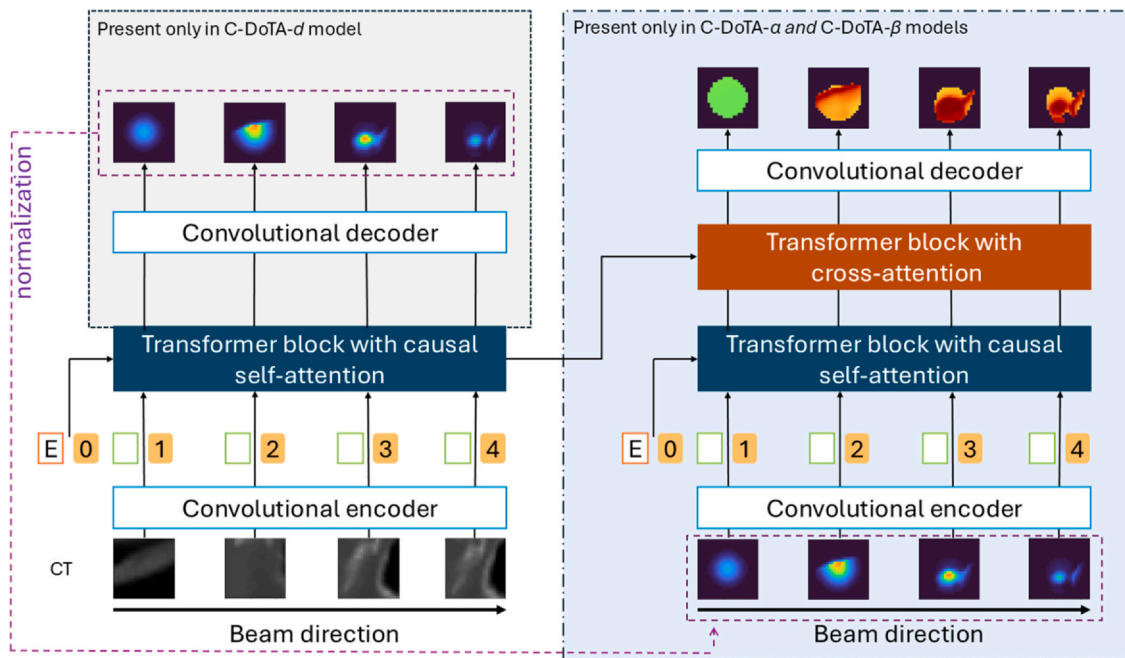


Figure 1. C-DoTA architectures. Left: Vanilla DoTA, with a convolutional encoder passing the normalized BEV-CT as a sequence of 2D images to the transformer block, together with the normalized energy E, used for dose prediction (C-DoTA-d). Right: The architecture for predicting α and β parameters is expanded by an additional convolutional encoder and transformer with the predicted dose as an input. The output of both transformers for CT and predicted dose is finally processed by a third transformer block using cross attention.

Model training details

The loss function used for the vanilla DoTA proved insufficient for our C-DoTA-model to learn the sharper gradients at the carbon Bragg peaks. The custom loss function was designed based on mean squared error (MSE) with 3 additional components, where y is the GT value and \tilde{y} is the predicted value:

Masked mean squared error:

$$\text{Masked MSE} = \frac{1}{n} \sum_{i: |y_i - \tilde{y}_i| > \epsilon^2} (y_i - \tilde{y}_i)^2 \quad (2.3)$$

where n is the cardinality of the set of voxels with a squared error above the threshold ϵ^2 (0.001 Gy for dose, 1/Gy for alpha and 1/Gy² for beta) at the given epoch. Unlike the traditional MSE loss, this masked loss excludes a high number of voxels with very small differences, forcing the model to pay more attention to the regions with high squared dose error, that is, emphasizing the narrow Bragg Peak where gradients are high.

Depth-wise partial gradient loss:

$$\text{Depth Loss} = \frac{1}{m} \sum_i^m \left(\frac{\partial y_i}{\partial z} - \frac{\partial \tilde{y}_i}{\partial z} \right)^2 \quad (2.4)$$

where the partial derivatives of the true and predicted responses with respect to depth z for each voxel i are calculated as

$$\frac{\partial y_i}{\partial z} = \frac{y_{j:z_j=z_i+\Delta z} - y_i}{\Delta z} \quad (2.5)$$

$$\frac{\partial \tilde{y}_i}{\partial z} = \frac{\tilde{y}_{j:z_j=z_i+\Delta z} - \tilde{y}_i}{\Delta z}, \quad (2.6)$$

with j being the index of the voxel in the next depth layer adjacent to the voxel i , Δz is the depth spacing and m is the total number of voxels.

Region masked mean squared error:

$$\text{Region masked MSE} = \frac{1}{l} \sum_{i: y_i > \theta} (y_i - \tilde{y}_i)^2 \quad (2.7)$$

where l is the cardinality of the set of voxels where GT values are above the threshold θ (60% of the sample maximum) at the given epoch. The loss increases the weight of the error in the regions with high values. These regions have the highest error if the model is trained on MSE only. The decision of the threshold was based on an empirical study.

The final loss is defined as the weighted sum of the separate MSE losses with the weight factors listed in [Table S1](#):

$$\text{Custom loss} = \text{MSE} + w_{\text{Mask}} \cdot \text{Masked MSE} + w_{\text{Depth}} \cdot \text{Depth Loss MSE} + w_{\text{Region}} \cdot \text{Region masked MSE} \quad (2.8)$$

To find the optimal weight, we trained models with different sets of weights in the interval [0,1] and a step of 0.2. The optimal weight values were chosen based on the evaluation of the validation dataset. Not all the weights were used in all models.

The C-DoTA models were first trained for two times 28 epochs using only a MSE loss, similar to the vanilla model. This was followed by two more times 28 epochs using a custom loss. The learning rate was rescheduled after each block of 28 epochs to prevent convergence in a local minimum.

Table 1
GPR 1%/1 mm results.

Parameter	Value range	GPR (1%/1 mm) [%]				
		Median	Min	10%-tile	90%-tile	Max
D	0-2.6 Gy	99.76	92.06	98.70	99.96	100
α	0-1.6 Gy ⁻¹	99.14	86.36	96.72	99.87	100
β	0-0.03 Gy ⁻²	98.74	85.26	96.71	99.45	99.87

AI dose mask was applied to both GT and AI output.

Evaluation

As evaluation metrics, we used the GPR with a 1%/1 mm criteria and a 10% threshold. The 10% threshold was applied to mitigate the influence of the metric in low-dose regions, as low values artificially inflated the GPR. In the results, the GPR was computed only for voxels where the C-DoTA-d prediction exceeded 1% of its maximum, that is, within the AI dose mask. This was necessary due to the impact of the AI dose mask on the alpha and beta predictions. The GPR based on the MC mask for the MC label can be found in [Supplementary Materials](#).

Furthermore, the consistency between prediction and GT was checked with dice coefficients²⁷ by calculating iso-dose contours for C-DoTA-d at 1%, 30%, and 70% of the sample maximum. The maximum of the GT sample was used to calculate relative metrics such as MSE. To estimate the impact of voxel exclusion due to differences in the 1% dose level between AI and MC, we counted voxels present in the MC mask but not in the AI dose mask and vice versa.

In addition to 3D-dose analyses, integrated depth-dose (IDD) profiles were assessed. Unlike GPR for IDD metrics, the MC mask was applied to GT, and the AI dose mask to the AI output to calculate the carbon range R80 of the Bragg peak, relative dose difference (ΔDmax) and MSE.

To automatically compute the carbon range R80 of the Bragg peak, the Python Shapely package was used to find the interpolated intersection value at 80% of the sample maximum. This allowed us to evaluate the range accuracy of the model dose output.

To estimate the model uncertainty distribution, we applied Monte Carlo dropout (MCD).²⁸ MCD is an efficient method for estimating the stability of a model's performance through multiple forward inferences with an activated dropout layer, thereby simulating multiple trained models. The evaluation was performed with a dropout rate of 0.2 on the complete prediction chain (and the dropout layers for both training and inference were the same as in the vanilla DoTA model).

In our setup, the alpha and beta predictions are strongly correlated to the quality of the C-DoTA-d dose input. In this work, we evaluated the propagated uncertainty using the MCD C-DoTA-d outputs performed with activated dropout layers. It means that the output of C-DoTA-d is used as input for C-DoTA- α and C-DoTA- β models. After 30 runs, the average value and standard deviation per sample were analyzed. For the graphic representation of the MCD-based uncertainty, a standard dose prediction was used as input to show the uncertainty regions of the models only. Quantitative statements on model uncertainty were obtained by evaluating the gamma passing rate of the 30 individual predictions.

Results

[Table 1](#) shows the results of GPR 1%/1 mm with the 10% threshold applied to the output of all 3 C-DoTA models (absorbed dose (D), alpha (α) and beta (β)). The median value for all 3 outputs is around 99%. 5%-tile values are 98.09% for absorbed dose, 95.60% for alpha and 95.73% for beta, which shows that most of the data is close to the expected values. The results for the MC mask are presented in [Supplementary Materials \(Table S2\)](#).

The top of [Figure 2](#) shows the output of the best absorbed dose prediction with GPR 1%/1 mm of 100% as well as the corresponding

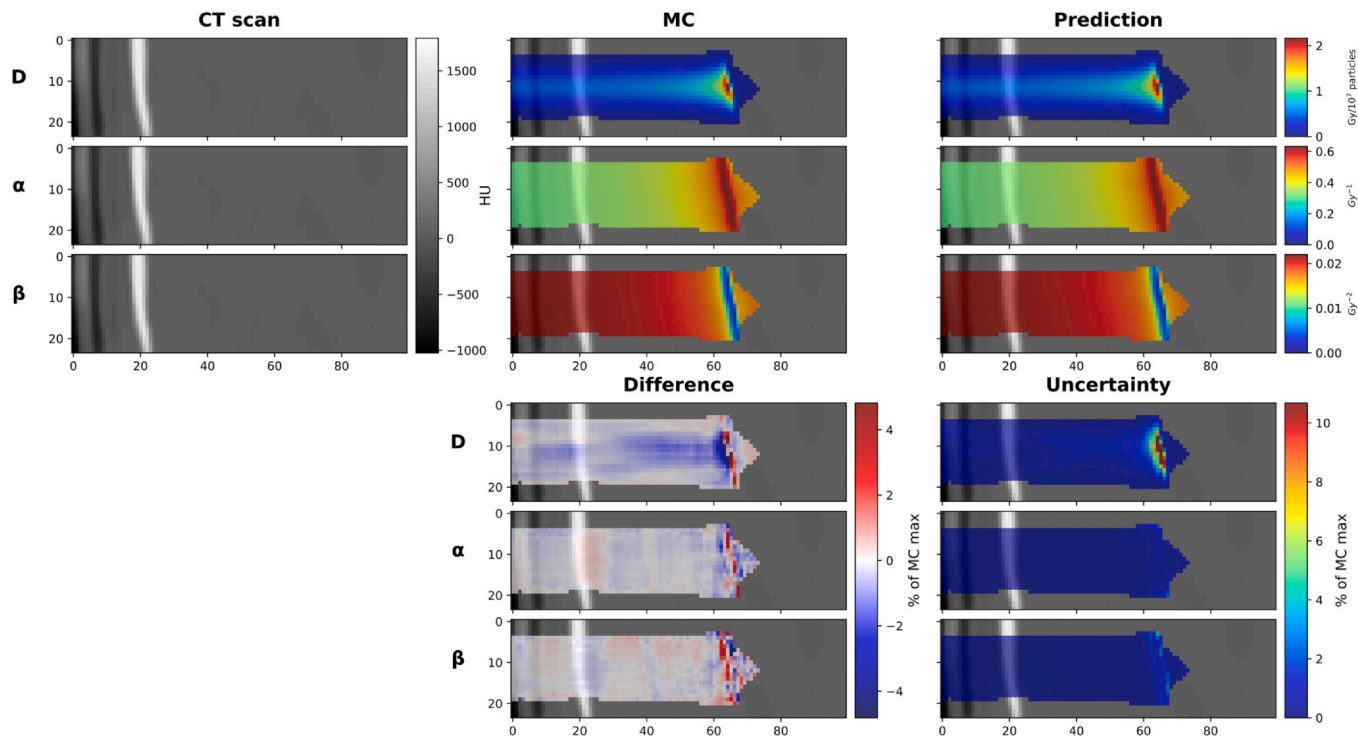


Figure 2. Best result in terms of the GPR for the physical dose. The GPR with a 1%/1 mm was top row absorbed dose (D): 100%, middle row alpha (α): 100%, bottom row beta (β): 99.87%. The top part: on the left, a slice of the BEV CT is shown; in the center, the MC GT is shown; and the C-DoTA predictions on the right. The bottom part: on the left side, the difference (MC minus C-DoTA); on the right side, MCD uncertainty. For better visibility, the maximum and minimum values of the color scale in the difference plots were set to the maximum between the absolute values of the 99.9%-tile and the 0.1%-tile. X and Y axis of all images are in mm.

alpha (100%) and beta (99.87%) predictions. In contrast, the bottom one shows the difference between GT and the prediction, as well as the corresponding MCD uncertainty maps. The results with MC mask applied to GT is showed in Figure S4. The top of Figure 3 shows a test sample with the worst GPR for physical dose prediction, at 92.06%, with alpha at 89.16% and beta at 90.62%. The bottom part illustrates its difference and MCD uncertainty. The worst-case sample has a highly inhomogeneous structure, featuring bones and a complex nasal cavity, where the Bragg peak is positioned exactly on the boundary between the regions of high and low density, resulting in a complex shape of the beamlet. The main uncertainty region for the dose is around the Bragg peak region for all the models, where the gradient is sharp. This somewhat correlates with the difference in the best-case scenario; however, no such relation is seen in the worst-case scenario. The results with MC mask applied to GT is showed in Figure S5.

The worst case of alpha (Figure S6) and beta (Figure S7) predictions can be found in Supplementary Materials.

To illustrate the effect of the custom loss, the results of the GPR 1%/1 mm for the models trained for 56 and 112 epochs on MSE can be found in Supplementary Materials (Table S3 and Table S4). While the physical dose prediction does not benefit much in GPR, the alpha and beta have gained $\approx 18\%$ in minimum values after 112 epochs. The additional 56 epochs improved the GPR by further $\approx 3\%$ at minimum.

The result of the MCD uncertainty prediction is summarized in Table 2 in the form of the metrics for the GPRs evaluated for each of the individual 30 MCD predictions. The median of the mean standard deviation was below 0.5% for all the models, which shows high reliability of the prediction.

Table 3 presents a more detailed analysis of the test dataset. The mode of the range difference is almost 0 mm, while the automatic detection of R80 causes some extreme differences. The median relative difference in the dose maximum is 0.9%. The root mean square error (RMSE) indicates that the average difference between the GT and

prediction is 0.63%, or cumulatively, 0.62 Gy/10⁷ primaries per 1 mm depth.

Examples of the extremes are shown in Figure S8. The left plot depicts the maximum difference in R80, while the middle one shows the minimum difference. It is clearly visible that the predicted IDD profile closely resembles the MC one, but due to its complex shape, the R80 value obtained is highly sensitive to minor changes, in this case causing a shift between the first and second peak in the IDD profile. In the right plot, the dose loss in the Bragg peak region is caused by a dense, bony CT structure, which also poses a challenge for PBAs to calculate. The corresponding C-DoTA outputs can be found in Figure S9.

The evaluation of the Dice coefficient between MC and C-DoTA predictions indicates a high level of agreement between the isodose lines. A small number of voxels primarily caused low values in some samples, where the relative impact of a small error is greater. The 10 and 90 percentile values are close to the median, demonstrating the stability of the C-DoTA results.

There is also a fraction of voxels that, due to the threshold, are present in GT but not in the prediction, and vice versa (Table S5). This fraction of the lost voxels due to thresholding is 8% with mean of 0.02 Gy/10⁷ primaries and a standard deviation of 0.003 Gy/10⁷ primaries.

As an additional test, not included in the tables, we checked the behavior of our models on some extreme cases, such as a lot of air in the nasal cavities due to a possible surgery before the therapy or for beam paths that pass many bone-tissue interfaces. This data is presented in Supplementary Material to this manuscript (Figure S10).

Finally, we analyzed the speed of the C-DoTA dose engine versus the FRED GPU MC. The time measurements were performed on a GPU Node with 8 AMD Radeon Instinct Mi100 32 GB video cards, 96 CPUs, and 384 GB RAM. Only 1 out of 8 GPUs was used, and the whole node was still blocked for other users to obtain a clean measurement of the runtime. The time for C-DoTA, including loading the samples into the model and dose pre-processing for alpha and beta models (but not the

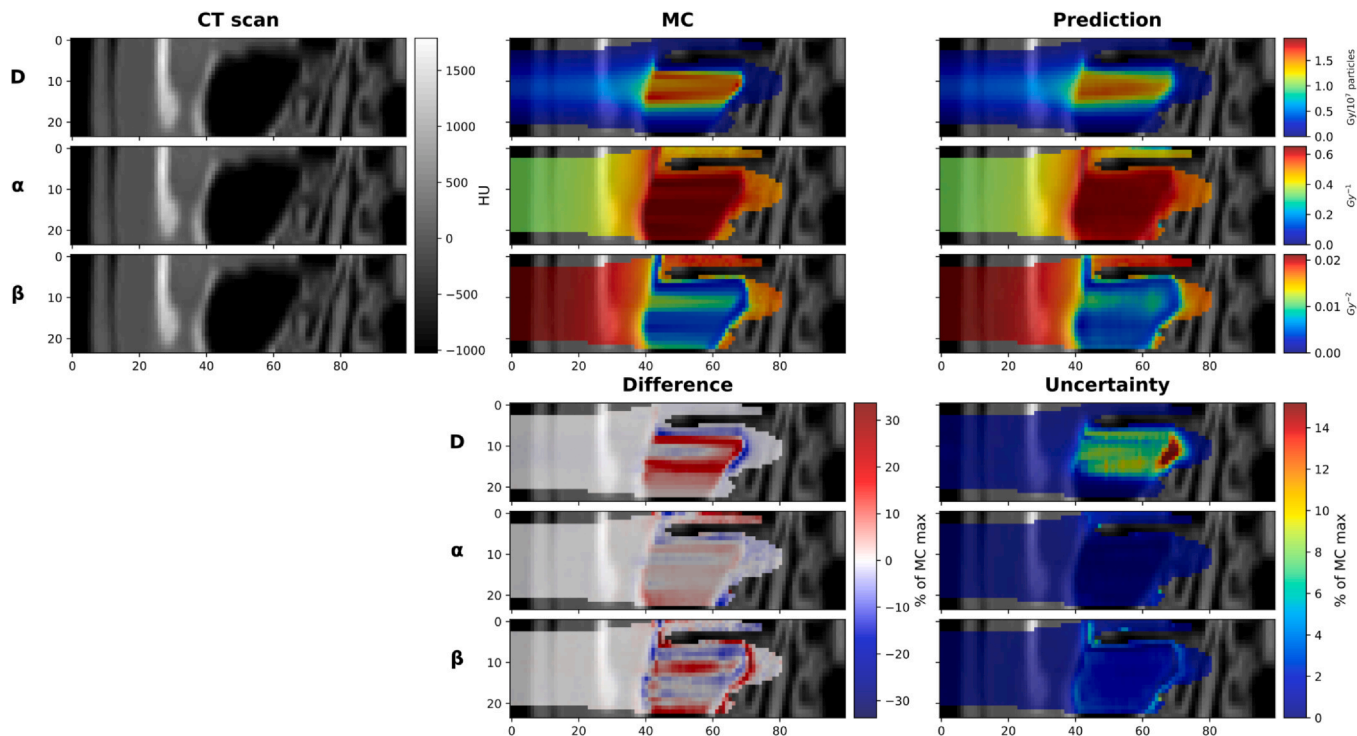


Figure 3. Worst result in terms of the GPR for the physical dose. The GPR with a 1%/1 mm was top row absorbed dose (D): 92.06%, middle row alpha (α): 89.16%, bottom row beta (β): 90.62%. The top part: on the left, a slice of the BEV CT is shown; in the center, the MC GT is shown; and the C-DoTA predictions on the right. The bottom part: on the left side, the difference (MC minus C-DoTA); on the right side, MCD uncertainty. For better visibility, the maximum and minimum values of the color scale in the difference plots were set to the maximum between the absolute values of the 99.9%-tile and the 0.1%-tile. X and Y axis of all images are in mm.

interpolation for obtaining the BEV CT), was 0.032 ± 0.001 seconds (0.007 ± 0.0002 seconds for the absorbed dose only). The samples were randomly selected from the test dataset. C-DoTA is capable of predicting > 400 times faster than MC FRED (14 seconds) for our setup.

Discussion

This work presents the first AI-based dose engine for carbon ion therapy capable of predicting both absorbed dose and the α and β parameters required for RBE-weighted dose calculation and optimization. The high GPRs (> 98% median) and strong spatial agreement (Dice coefficient 0.95 at 1% isodose) demonstrate that the proposed framework delivers robust predictions across a wide range of anatomical configurations. Accuracy loss in highly heterogeneous regions—also a challenge for PBAs—was modest and primarily affected the high-gradient dose falloff. These cases highlight the importance of targeted data augmentation or hybrid MC–AI workflows to further improve performance in extreme anatomies, such as post-surgical cavities or dense bony structures.

Table 2

Mean (top row) and standard deviation (bottom row) of GPR per sample after 30 runs with MCD of 0.2.

Parameter	Value range	GPR (1%/1 mm) [%]				
		Median	Min	10%-tile	90%-tile	Max
D	0-2.6 Gy	99.59	92.62	96.62	99.89	100
		0.18	0.01	0.06	0.67	3.49
α	0-1.6 Gy ⁻¹	98.62	86.97	96.42	99.57	99.96
		0.47	0.04	0.15	1.52	6.28
β	0-0.03 Gy ⁻²	98.20	85.59	96.12	99.05	99.52
		0.39	0.05	0.16	1.42	9.57

AI dose mask was applied to both GT and AI output.

Clinically, the model’s independence from simulation statistics, beam energy, and anatomy for inference time makes it well-suited for online adaptive planning and 4D robust optimization. Parallelization across GPUs could potentially reduce runtime to milliseconds for full plans, enabling real-time replanning during treatment sessions. The high-quality PB data output of the model can be directly used for optimization, as opposed to other MC²⁹ or AI models³⁰ that compute or predict the full dose of a plan.

However, several limitations must be acknowledged. The training data were restricted to head-and-neck patients, a fixed horizontal beamline, and a therapeutic energy range of 115-260 MeV/u. A broader generalization will require expanding the dataset to include more anatomical sites, beam geometries, and beamlines. Additionally, while α and β predictions matched MC values, their biological accuracy depends on the LEM and requires further validation against experimental radiobiology data. Future research should address the following issues: the effect of the absence of the low-dose region in the C-DoTA prediction (8% of voxels on average) and the error introduced by alpha and beta predictions on RBE-weighted dose, both of which need to be investigated in the context of a full plan. This full plan can

Table 3
Additional evaluation metrics.

		Median	Min	10%-tile	90%-tile	Max
IDD Curve	$\Delta R80$ mm	-0.02	-24.48	-0.44	0.32	26.20
	ΔD_{max} %	0.90	-11.51	-2.25	4.32	18.38
	RMSE %*	0.97	0.25	0.65	1.61	6.12
FRED vs C-Dota	RMSE %*	0.14	0.05	0.09	0.23	0.63
	Dice (1% isodose)	0.95	0.84	0.93	0.96	0.98
	(n voxels in GT)	(19 241)	(11 531)	(14 455.5)	(24 566)	(38 586)
	Dice (30% isodose)	0.90	0.29	0.77	0.96	0.99
	(n voxels in GT)	(640)	(301)	(418)	(2016.5)	(5907)
	Dice (70% isodose)	0.83	0.16	0.68	0.92	1
(n voxels in GT)	(39)	(15)	(29)	(77)	(499)	

* values calculated using the AI dose mask.

also be compared against another MC model, featuring a more complete physics implementation at longer calculation times, to benchmark the simplifications inherent to FRED's GPU implementation. Moreover, the consistency between the base data for optimization and dose calculation reduces the error of the final plans. Finally, to generalize the model for various treatment sites, the inclusion of tissue-specific alpha/beta ratios needs to be studied. FRED is a GPU-based MC tool for fast simulation and includes accurate representation of the beam transport and interaction at the voxel level, delivering the most important characteristics needed to train the C-DoTA model. Nevertheless, it is not a clinically tested tool and does not consider refragmentation. For clinical application the models need to be fine-tuned on the data created with more accurate and currently slower MC tools.

Overall, this proof-of-concept demonstrates that deep learning can deliver input parameters for the biologically-weighted dose needed for carbon ion therapy with clinical runtimes and MC-quality, paving the way for integration into routine workflows and wider adoption of this modality.

Conclusions

We developed an AI model able to calculate the input parameters for RBE-weighted dose calculation in carbon-ion therapy with MC accuracy but at 400-fold faster computation time. The model's speed can potentially enable rapid online re-planning, which is particularly valuable for addressing inter-fractional anatomical changes. By removing the need for MC, C-DoTA will improve carbon ion therapy accuracy, enabling safer and more effective treatments. Future directions include lowering the prediction threshold to 0.5% (clinically required); collaboration with C-ion clinical centers to test robustness across beamlines and patient cohorts; and embedding the model into clinical workflows (eg, TRiP98) to evaluate end-to-end plan quality.

CRedit authorship contribution statement

A. Quarz: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Validation, Visualization, Writing – original draft. **A. De Gregorio:** Software, Data curation, Writing – review and editing. **G. Franciosini:** Software, Data curation, Writing – review and editing. **A. Schiavi:** Software, Writing – review and editing. **Z. Perko:** Conceptualization, Methodology, Supervision, Writing – review and editing. **L. Volz:** Conceptualization, Methodology, Supervision, Writing – review and editing. **C. Hoog Antink:** Supervision, Writing – review and editing. **V. Patera:** Supervision, Writing – review and editing. **M. Durante:** Conceptualization, Project administration, Writing – review and editing. **C. Graeff:** Conceptualization, Software, Supervision, Project administration, Funding acquisition, Writing – review and editing.

Declaration of Conflicts of Interest

The authors declare the following financial interests/personal relationships, which may be considered as potential competing interests:

Zoltan Perko reports a relationship with Radformation Inc that includes: employment. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The publication is funded by the Open Access Publishing Fund of GSI Helmholtzzentrum fuer Schwerionenforschung.

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at [doi:10.1016/j.ijpt.2026.101309](https://doi.org/10.1016/j.ijpt.2026.101309).

References

- Durante M, Orecchia R, Loeffler JS. Charged-particle therapy in cancer: clinical uses and future perspectives. *Nat Rev Clin Oncol*. 2017;14:483–495. <https://doi.org/10.1038/nrclinonc.2017.30>
- Durante M, Paganetti H. Nuclear physics in particle therapy: a review. *Rep Prog Phys*. 2016;79:096702. <https://doi.org/10.1088/0034-4885/79/9/096702>
- Dedes G, Parodi K. Monte Carlo simulations of particle interactions with tissue in carbon ion therapy. *Int J Part Ther*. 2015;2:447–458. <https://doi.org/10.14338/IJPT-15-00021>
- Simoni MD, Fischetti M, Gioscio E, et al. FRED: a fast Monte Carlo code on GPU for quality control in particle therapy. *J Phys Conf Ser*. 2020;1548:012020. <https://doi.org/10.1088/1742-6596/1548/1/012020>
- Mein S, Choi K, Kopp B, et al. Fast robust dose calculation on GPU for high-precision 1H, 4He, 12C and 16O ion therapy: the FRoG platform. *Sci Rep*. 2018;8:14829. <https://doi.org/10.1038/s41598-018-33194-4>
- Janson M, Glimelius L, Fredriksson A, Traneus E, Engwall E. Treatment planning of scanned proton beams in RayStation. *Med Dosim*. 2024;49:2–12. <https://doi.org/10.1016/j.meddos.2023.10.009>
- Wolf M, Anderle K, Durante M, Graeff C. Robust treatment planning with 4D intensity modulated carbon ion therapy for multiple targets in stage IV non-small cell lung cancer. *Phys Med Biol*. 2020;65:215012. <https://doi.org/10.1088/1361-6560/aba1a3>
- Paganetti H, Botas P, Sharp GC, Winey B. Adaptive proton therapy. *Phys Med Biol*. 2021;66:22TR01. <https://doi.org/10.1088/1361-6560/ac344f>
- Graeff C, Volz L, Durante M. Emerging technologies for cancer therapy using accelerated particles. *Prog Part Nucl Phys*. 2023;131:104046. <https://doi.org/10.1016/j.pnpnp.2023.104046>
- Tan J, Parisi A, Tan KM, et al. Pioneering MCF MKM RBE-weighted dose calculation for carbon ion radiotherapy: development of a pencil beam algorithm and validation against Monte Carlo simulation on clinical CT data. *Phys Med Biol*. 2025;70:155009. <https://doi.org/10.1088/1361-6560/adf36f>
- Krämer M, Scholz M. Rapid calculation of biological effects in ion radiotherapy. *Phys Med Biol*. 2006;51:1959–1970. <https://doi.org/10.1088/0031-9155/51/8/001>
- Ruangchan S, Palmans H, Knäusel B, Georg D, Clausen M. Dose calculation accuracy in particle therapy: comparing carbon ions with protons. *Med Phys*. 2021;48:7333–7345. <https://doi.org/10.1002/mp.15209>
- Deng Y, Chen Z, Wang Q, et al. Comparison of pencil beam and Monte Carlo calculations with ion chamber array measurements for patient-specific quality assurance. *Radiat Med Prot*. 2022;3:115–122. <https://doi.org/10.1016/j.radmp.2022.07.002>
- Zhang H, Li Q, Liu X, Ma Y, He P. Physica Medica validation and testing of a novel pencil-beam model derived from Monte Carlo simulations in carbon-ion treatment planning for different scenarios. *Phys Med*. 2022;99:1–9. <https://doi.org/10.1016/j.ejmp.2022.04.018>
- Pastor-Serrano O, Perko Z. Millisecond speed deep learning based proton dose calculation with Monte Carlo accuracy. *Phys Med Biol*. 2022;67:105006. <https://doi.org/10.1088/1361-6560/ac692e>

16. Radonic D, Xiao F, Wahl N, et al. Proton dose calculation with LSTM networks in presence of a magnetic field. *Phys Med Biol.* 2024;69. <https://doi.org/10.1088/1361-6560/ad7f1e>
17. Neishabouri A, Wahl N, Mairani A, Köthe U, Bangert M. Long short-term memory networks for proton dose calculation in highly heterogeneous tissues. *Med Phys.* 2021;48. <https://doi.org/10.1002/mp.14658>
18. Zhang X, Hu Z, Zhang G, Zhuang Y, Wang Y, Peng H. Dose calculation in proton therapy using a discovery cross-domain generative adversarial network (DiscoGAN). *Med Phys.* 2021;48. <https://doi.org/10.1002/mp.14781>
19. Zhang L, Holmes JM, Liu Z, et al. Beam mask and sliding window-facilitated deep learning-based accurate and efficient dose prediction for pencil beam scanning proton therapy. *Med Phys.* 2024;51:1484–1498. <https://doi.org/10.1002/mp.16758>
20. Wu C, Nguyen D, Xing Y, et al. Improving proton dose calculation accuracy by using deep learning. *Mach Learn Sci Technol.* 2021;2(1):015017. <https://doi.org/10.1088/2632-2153/abb6d5>
21. Xing Y, Zhang Y, Nguyen D, Lin MH, Lu W, Jiang S. Boosting radiotherapy dose calculation accuracy with deep learning. *J Appl Clin Med Phys.* 2020;21(8):149–159. <https://doi.org/10.1002/acm2.12937>
22. He R, Zhang H, Wang J, et al. Deep learning-based prediction of Monte Carlo dose distribution for heavy ion therapy. *Phys Imag Radiat Oncol.* 2025;34:100735. <https://doi.org/10.1016/j.phro.2025.100735>
23. De Simoni M, Battistoni G, De Gregorio A, et al. A data-driven fragmentation model for carbon therapy GPU-accelerated Monte-Carlo dose recalculation. *Front Oncol.* 2022;12:780784. <https://doi.org/10.3389/fonc.2022.780784>
24. Krämer M, Scholz M. Treatment planning for heavy-ion radiotherapy: calculation and optimization of biologically effective dose. *Phys Med Biol.* 2000;45:3319–3330. <https://doi.org/10.1088/0031-9155/45/11/314>
25. Schiavi A, Senzacqua M, Pioli S, et al. Fred: a GPU-accelerated fast-Monte Carlo code for rapid treatment plan recalculation in ion beam therapy. *Phys Med Biol.* 2017;62:7482–7504. <https://doi.org/10.1088/1361-6560/aa8134>
26. Elsässer T, Krämer M, Scholz M. Accuracy of the local effect model for the prediction of biologic effects of carbon ion beams in vitro and in vivo. *Int J Radiat Oncol Biol Phys.* 2008;71:866–872. <https://doi.org/10.1016/j.ijrobp.2008.02.037>
27. Dice LR. Measures of the amount of ecologic association between species. *Ecol Soc Am.* 1945;26:297–302.
28. Gal Y, Ghahramani Z. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. 33rd International Conference on Machine Learning, ICML 2016 2016;3:1651–1660.
29. Basaglia T, Bell ZW, D'Agostino D, et al. Geant4: a game changer in high energy physics and related applicative fields. *Future Gener Comput Syst.* 2024;159. <https://doi.org/10.1016/j.future.2024.05.042>
30. Neishabouri A, Bauer J, Abdollahi A, Debus J, Mairani A. Real-time adaptive proton therapy: an AI-based spatio-temporal mono-energetic dose calculation model (CC-LSTM). *Comput Biol Med.* 2025;188. <https://doi.org/10.1016/j.compbiomed.2025.109777>