

4D Pencil Beam Treatment Plan with Conditional Weight Predictions

Nair N von Mühlenen¹[0009–0002–5122–5480], Florentin Bieder¹[0000–0001–9558–0623], and Philippe C Cattin¹[0000–0001–8785–2713]

Department of Biomedical Engineering, University of Basel, 4001 Basel, Switzerland
nair.vonmuehlenen@unibas.ch

Abstract. *Objective.* This work aims to accelerate our four-dimensional (4D) proton pencil beam delivery strategy, which incorporates respiratory motion into a dynamic treatment plan, to improve dose conformity and treatment efficiency for gantry-less and magnet-free scanner designs. *Approach.* To accelerate the generation and adaptation of 4D treatment plans, we propose a conditional hybrid ResNet18-Transformer model for predicting pencil beam weights. The model predicts either the full set of pencil beam weights or a subset thereof. *Main Results.* The conditional model can successfully predict the remaining beam weights of a treatment, and the results suggest that extending this approach to full treatment prediction is feasible but requires further investigation. *Significance.* The acceleration of 4D treatment generation via pencil beam weight prediction takes us one step closer to the feasibility of treating mobile targets with simplified, gantry-less, and magnet-free scanner designs. Reducing system complexity while preserving dosimetric conformity may offer a pathway toward more accessible and cost-effective proton beam therapy for motion-affected tumours. Our code is available on github.com/NairVonMuehlenen/Pencil-Beam-Weight-Prediction.

Keywords: Proton therapy · 4D treatment planning · Weight prediction

1 Introduction

Proton therapy (PT) is an advanced form of radiation therapy used in the treatment of cancer [13]. Unlike conventional photon-based radiotherapy, PT uses protons to deliver radiation. This treatment benefits from the unique physical properties of protons, particularly the Bragg Peak, which enable highly localised radiation deposition [12]. Beyond this peak, the radiation dose drops rapidly, so virtually no radiation is delivered beyond the target depth. This characteristic dose distribution offers significant advantages for cancer treatment as it can minimise radiation exposure to surrounding healthy tissue, organs at risk (OAR) and other sensitive structures [8,9]. Despite this significant advantage, PT also presents several challenges. The treatment is highly sensitive to patient and tumour motion, which can affect the placement of the delivered dose. In addition, the infrastructure required for PT is complex and expensive, limiting its availability to specialised centres and making it inaccessible to many worldwide.

This work of von Mühlenen et al. [10] introduced a new 4D planning framework for mobile tumours and its possible application in gantry-less and even magnet-less PT systems. It showed promising future applications of more cost-effective gantry-less PT systems. The system takes as input a CT scan, the planning target volume (PTV), and the corresponding respiratory motion fields, and outputs the pencil beam positions, their delivery times relative to the start of treatment, and their weights. The delivery time includes a nominal dwell time and the travel time between spots. The beam weights are optimised to achieve homogeneous dose coverage of the target while minimising radiation exposure to healthy tissue. The optimisation minimises an objective function composed of common PT metrics such as $V_{95\%}$ and $D_{95\%}$. With the optimised pencil beam weights, the treatment plan can be adapted. In particular, removing spots with negligible weights and adapting the nominal dwell time to the actual dwell time, altering the temporal structure of the delivery sequence. These properties motivate an iterative planning strategy, in which beam weights are recomputed to refine and potentially simplify the delivery plan. However, conventional weight optimisation is too slow for such an application. This motivated the development of a fast weight-prediction model that enables rapid recomputation and adjustment of treatment plans while maintaining treatment-plan quality. An overview of the proposed architecture is shown in Fig. 1.




Fig. 1: Schematic diagram of the conditional hybrid ResNet18-Transformer: It consists of a ResNet18 encoder and a Transformer decoder. The inputs to the model are the motion-deformed CT scan with the current accumulated dose distribution, the prescribed Gray (PG) and a predetermined sequence of spots.

1.1 Related Work

Recent years have seen significant progress in 4D treatment planning for PT. Modern approaches now enable robust, patient-specific dose calculations that account for respiratory motion and anatomical changes, improving treatment accuracy and safety [5,7]. Several centres have implemented clinically feasible, custom 4D dose reconstruction and accumulation tools within commercial planning systems, allowing for retrospective and prospective plan evaluation and supporting adaptive plan evaluation with processing times of approximately 10 minutes per treatment fraction and only limited manual intervention [4,7]. Four-dimensional robust optimisation (4DRO) strategies have demonstrated improved target coverage and reduced interplay effects compared to traditional 3D planning [5,6]. Adaptive planning methods, such as dose-mimicking and template-based approaches, are being explored to automate plan adaptation during treatment, thereby further enhancing robustness and sparing OAR [2,4].

Despite these advances, several challenges remain. (1) Computational Demands: Full 4D robust optimisation remains computationally intensive, though strategies to reduce the number of phases or use surrogate phases have been proposed to balance speed and quality [2,6]. (2) Motion and Anatomy Variability: Accurate modelling of patient-specific motion and anatomical changes over the treatment course is complex and requires repeated imaging or advanced motion modelling [4,7]. (3) Clinical Integration: While proof-of-concept and early clinical implementations exist, large-scale, fully automated, and prospective 4D adaptive workflows are not yet standard in most clinics [2,7].

2 Method

Since weight optimisation is the most computationally expensive step in the 4D treatment planning process, we propose a significantly faster weight-prediction model for treatment adaptation and, potentially, for full treatment beam weight prediction. Specifically, we employ a conditional hybrid ResNet18-Transformer model that predicts either the full set of pencil beam weights or a subset thereof. To ensure comparable plan quality, the model is trained using the same metric-based loss function as the optimisation framework shown in von Mühlenen et al. [10], promoting homogeneous target coverage while minimising dose to healthy tissue.

2.1 Dataset

In this study, we use simulated data representing a moving liver tumour. This choice provides full control over the entire workflow and ensures that all components can be evaluated under controlled, reproducible conditions. The data generation pipeline follows the aforementioned work [10], with the dataset extended from spherical targets to more organically shaped tumour geometries. Consequently, the tumours vary in shape, spatial extent, and location between

phantoms rather than being defined by a single fixed diameter. For each geometry, corresponding optimised pencil-beam weights are available as reference solutions. The simulated phantom comprises different tissue-equivalent regions with varying Hounsfield units, densities, and thicknesses. The phantoms are represented on a $100 \times 100 \times 100$ voxel grid with a spatial resolution of $1 \times 1 \times 1 \text{ mm}^3$. Figure 2 shows a CT slice from the simulated data on the right and the corresponding dose distribution obtained using the optimised beam weights on the left. For each phantom, we compute a set of planned pencil beam positions and possible delivery times. The complete data-generation and treatment-planning pipeline, is publicly available in our GitHub repository.




Fig. 2: Example of a simulated phantom: On the right, a CT centre slice showing heterogeneous tissue-equivalent regions. On the left, the corresponding radiation map of the centre slice with optimised weights.

2.2 Model Architecture

The proposed conditional hybrid ResNet18–Transformer model consists of a ResNet18 encoder [1] and a Transformer decoder [14]. The encoder receives the motion-deformed CT scan along with the current accumulated dose distribution, which represents the dose delivered up to the current point in time. The extracted feature representation is concatenated with the prescribed dose (PG) and provided to the Transformer decoder. The decoder receives, as queries, a predetermined sequence of Fourier-embedded beam positions and delivery times, $(x_i, y_i, z_i, t_i)_{i=1}^N$. Depending on the application, the query sequence may include either the full set of planned spots or only the remaining delivery sequence from a given time onward. In the latter case, the accumulated dose supplied to the encoder corresponds to the dose delivered up to the first queried beam. The decoder predicts one weight per queried pencil beam. A schematic overview of the proposed architecture is shown in Fig. 1.

2.3 Training Objective

Rather than directly regressing the optimised beam weights, the proposed network is trained to minimise the objective function used during treatment plan optimisation. From the predicted beam weights, we calculate the dose distribution within the phantom using the analytical model of Schaffner et al. [11]. The resulting distribution is then evaluated using the same objective function employed in the optimisation framework von Mühlenen et al. [10]. This ensures that training directly optimises treatment plan quality instead of agreement with a particular set of beam weights. The loss is formulated in terms of the dosimetric metrics $V_{95\%}$, $D_{95\%}$, $D_{max}TV$ and $D_{max}HT$ and is defined as

$$\mathcal{L} = \mathcal{L}_V + \mathcal{L}_D + c_{max} (\mathcal{L}_{D_{max}TV} + \mathcal{L}_{D_{max}HT}), \quad (1)$$

where c_{max} controls the relative contribution of the maximum-dose penalties. The individual loss terms are given in Equations 2 - 5, while the corresponding optimisation metrics are summarised in Table 1. Each loss component is computed relative to the prescribed physical dose.

$$\mathcal{L}_V = \max\{L_V - V_{95\%}, 0\} \quad (2)$$

$$\mathcal{L}_D = \max\{L_D - D_{95\%}, 0\} \quad (3)$$

$$\mathcal{L}_{D_{max}TV} = \max\{D_{max}TV - L_{max}TV, 0\} \quad (4)$$

$$\mathcal{L}_{D_{max}HT} = D_{max}HT - L_{max}HT \quad (5)$$

Table 1: Used optimisation metrics depended on the prescribed Gray (PG).

Notations	Definition	Optimisation Metric
L_V	The ideal value for $V_{95\%}$	98 % of PG
L_D	The ideal value for $D_{95\%}$	PG
$L_{max}TV$	The maximum for $D_{max}TV$	< 120 % of PG
$L_{max}HT$	The maximum for $D_{max}HT$	< 70 % of PG

2.4 Training Procedure

The network was trained using the Adam optimiser [3] with a learning rate of 10^{-5} and a batch size of 2, resulting in approximately 8,800 optimisation steps per epoch. The model was trained for up to 70 epochs, with the checkpoint achieving the lowest validation loss selected for evaluation. The dataset comprised 20 simulated phantom patients, with one patient held out for testing and the remaining 19 patients divided into five folds for training and validation. Dose augmentation was applied to generate prescribed physical doses ranging from 1 to 80 Gy. The loss weight c_{max} was set to 0.5. All experiments were conducted on a single NVIDIA A100 GPU.




Fig. 3: Dose maps for full- and partial-beam weight prediction. Columns, from left to right, correspond to the fully predicted treatment plan (100%), the remaining 75 %, 50 %, and 25 % of the treatment, and the ground truth. The first row shows the accumulated partial dose distribution used as input to the model; the second row shows the dose contribution from the predicted beam weights; and the third row shows the hybrid treatment plan obtained by combining the reference with the predicted beam weights.

3 Results

We evaluated the proposed model by predicting both the complete set of beam weights and the remaining beam weights at different stages of treatment delivery. For the evaluation, we utilised 10 Gy as the PG. Figure 3 compares the corresponding dose distributions for the fully predicted, partially predicted, and fully optimised ground truth (GT) treatment plans. The first row illustrates the accumulated dose distribution provided as input to the model, obtained from the GT plan. The second row presents the dose distributions resulting from the predicted beam weights, while the final row shows the combined distributions obtained by merging the optimised and predicted beam weights. These hybrid dose distributions are used to evaluate model performance using the loss defined in Equation 1. Partial predictions were evaluated after 25 %, 50 %, and 75 % of the treatment had been delivered, corresponding to predictions for the remaining 75 %, 50 %, and 25 % of the treatment, respectively. As expected, prediction accuracy improved as the proportion of beam weights requiring prediction decreased. The closest agreement with the fully optimised reference was achieved when only the final 25 % of the treatment required prediction, followed by prediction of the

remaining 50 % and 75 % of the treatment. Predicting the complete set of beam weights resulted in the largest deviations from the optimised reference.

Figure 4 presents the corresponding dose volume histograms (DVHs) for the dose distributions shown in Figure 3. Figures 4a-4c compare the fully optimised GT plan (lilac), the partially delivered GT plan provided as model input (blue), and the hybrid dose distribution obtained by combining the optimised and predicted beam weights (black). The DVH demonstrate that, in all partial prediction scenarios, the completed treatment plans are closer to the fully optimised reference than the partially delivered GT plans. This indicates that the predicted beam weights consistently improve the dosimetric quality of the partially delivered plans, irrespective of the prediction stage. Figure 4d compares the DVHs of the fully optimised GT plan and the treatment plan generated using the fully predicted beam weights.




Fig. 4: Dose-volume histograms (DVHs) for the treatment plans shown in Figure 3. Figures 4a-4c: The lilac curve represents the fully optimised reference plan, the blue curve the accumulated dose at the prediction stage, and the black curve the hybrid treatment plan obtained after incorporating the predicted beam weights. Figure 4d compares the DVHs of the fully optimised ground truth plan (lilac) and the treatment plan generated using the fully predicted beam weights (black).

4 Discussion

In this work, we introduced a conditional hybrid ResNet18–Transformer model for predicting pencil beam weights in 4D treatment planning. It can predict partial beam weights while maintaining the previously optimised dose distribution and could serve as an extension of the 4D treatment planning workflow for treatment plan adaptation. A key aspect of the proposed approach is the use of a physics-based loss function. Instead of minimising the difference between predicted and optimised beam weights, the network is trained through the resulting dose distribution using clinically relevant metrics. This allows the model to prioritise the final treatment objective, namely homogeneous target coverage, while limiting dose to healthy tissue, rather than learning a specific optimised solution. Consequently, the model is not required to replicate the exact optimisation process; instead, it learns to generate beam weights that satisfy the same constraints.

The Transformer architecture enables prediction of either the complete treatment plan or only the remaining beam weights from an intermediate stage of delivery. This capability is particularly relevant for iterative treatment planning, where beam weights may be updated after partial delivery to account for changes in the treatment state. The results demonstrate that the model can improve partially delivered treatment plans by predicting the remaining beam weights while retaining the previously delivered dose contribution.

The prediction accuracy improved as the number of beam weights to be predicted decreased. The best agreement with the fully optimised reference plans was achieved when predicting only the final fraction of the treatment delivery, while predicting the complete treatment plan resulted in larger deviations. This can be seen both in the dose distribution in Figure 3 and the corresponding DVHs in Figure 4. This behaviour is likely related to two factors. First, later delivered beams generally contribute less to the overall dose distribution, as the majority of the prescribed dose has already been delivered. Therefore, predicting these remaining beam weights represents a less challenging task. Second, the training distribution is inherently biased towards later treatment stages, as the model always receives queries corresponding to the final part of the delivery sequence. Whereas the earlier sections are queried much less frequently. Consequently, the network predicts the later stages of delivery more frequently than the early stages, which involve spots that contribute more significantly to the treatment.

Future work could investigate strategies to improve prediction performance for earlier treatment stages. One possible approach would be weighted sampling to increase the frequency of challenging early-stage prediction scenarios. Another approach would be to introduce intermediate training tasks, where the network predicts earlier beam weights while later beam weights remain fixed to their optimised values. Although this does not directly correspond to the temporal delivery process, it could allow the network to learn the more challenging initial beam weight predictions before being tasked with predicting the complete delivery sequence.

Several limitations of this study should be considered. First, the evaluation was performed exclusively on simulated phantom patient data. While this approach provides complete control over the 4D treatment-planning workflow and enables reproducible comparison with the optimisation framework, the dataset’s anatomical diversity remains limited. The current dataset contains similar structures and does not capture the full range of patient-specific variations encountered in clinical practice. Future work should therefore evaluate the proposed model on a larger and more anatomically diverse dataset, including clinical patient data where available. Second, the generalisability of the model to different scanners, delivery systems, and machine-specific parameters remains to be investigated. Since delivery times are determined by machine-specific scanner parameters, adaptation to other delivery systems may require additional training data or transfer learning.

Despite these limitations, a major advantage of the proposed approach is the potential substantial reduction in computational cost. While conventional beam weight optimisation requires several hours, the neural network prediction can be performed within milliseconds. This speed-up would enable repeated plan evaluation and refinement, which is impractical when relying solely on conventional optimisation. Therefore, fast beam weight prediction represents a promising direction towards iterative 4D treatment planning workflows, where plans can be continuously updated, simplified, and evaluated while maintaining a homogeneous dose coverage of the target and minimising radiation exposure to healthy tissue.

5 Conclusion

This work presented a conditional hybrid ResNet18–Transformer model for predicting pencil beam weights in 4D proton treatment planning. The Transformer architecture supports both complete and partial beam weight prediction, enabling efficient updates to treatment plans at arbitrary stages of delivery. Overall, this work demonstrates the feasibility of conditional, physics-informed beam weight prediction for partial treatment replanning. While prediction of complete treatment plans remains an open challenge, the proposed framework provides a promising foundation for future learning-based alternatives to conventional beam weight optimisation.

References

1. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
2. Kaushik, S., Stützer, K., Ödén, J., Fredriksson, A., Toma-Dasu, I.: Adaptive intensity modulated proton therapy using 4d robust planning: a proof-of-concept for the application of dose mimicking approach. *Physics in Medicine and Biology* **69** (2024). <https://doi.org/10.1088/1361-6560/ad75e0>
3. Kingma, D.P.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
4. Lebbink, F., Stocchiero, S., Fossati, P., Engwall, E., Georg, D., Stock, M., Knäusl, B.: Parameter based 4d dose calculations for proton therapy. *Physics and Imaging in Radiation Oncology* **27** (2023). <https://doi.org/10.1016/j.phro.2023.100473>
5. Liu, W., Schild, S., Chang, J.Y., Liao, Z., Chang, Y.H., Wen, Z., Shen, J., Stoker, J., Ding, X., Hu, Y., Sahoo, N., Herman, M., Vargas, C., Keole, S., Wong, W., Bues, M.: Exploratory study of 4d versus 3d robust optimization in intensity modulated proton therapy for lung cancer. *International journal of radiation oncology, biology, physics* **95** **1**, 523–33 (2016). <https://doi.org/10.1016/j.ijrobp.2015.11.002>
6. Mastella, E., Molinelli, S., Pella, A., Vai, A., Maestri, D., Vitolo, V., Baroni, G., Valvo, F., Ciocca, M.: 4d strategies for lung tumors treated with hypofractionated scanning proton beam therapy: Dosimetric impact and robustness to interplay effects. *Radiotherapy and oncology : journal of the European Society for Therapeutic Radiology and Oncology* **146**, 213–220 (2020). <https://doi.org/10.1016/j.radonc.2020.02.025>
7. Meijers, A., Jakobi, A., Jakobi, A., Stützer, K., Marmitt, G., Both, S., Langendijk, J., Richter, C., Knopf, A.: Log file-based dose reconstruction and accumulation for 4d adaptive pencil beam scanned proton therapy in a clinical treatment planning system: Implementation and proof-of-concept. *Medical Physics* **46**, 1140–1149 (2019). <https://doi.org/10.1002/mp.13371>
8. Mohan, R.: A review of proton therapy—current status and future directions. *Precision radiation oncology* **6**(2), 164–176 (2022)
9. Mohan, R., Grosshans, D.: Proton therapy—present and future. *Advanced drug delivery reviews* **109**, 26–44 (2017)
10. von Muehlenen, N.N., Bieder, F., Zhang, Y., Cattin, P.C.: Magnet-free proton therapy with 4d pencil beam delivery optimisation (2026), <https://arxiv.org/abs/2606.03562>
11. Schaffner, B., Pedroni, E., Lomax, A.: Dose calculation models for proton treatment planning using a dynamic beam delivery system: an attempt to include density heterogeneity effects in the analytical dose calculation. *Physics in Medicine & Biology* **44**(1), 27 (1999)
12. Schulz-Ertner, D., Tsujii, H.: Particle radiation therapy using proton and heavier ion beams. *Journal of clinical oncology* **25**(8), 953–964 (2007)
13. Smith, A.R.: Proton therapy. *Physics in Medicine & Biology* **51**(13), R491 (2006)
14. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)