

Article

Fast, Efficient and Flexible Particle Accelerator Optimisation Using Densely Connected and Invertible Neural Networks

Renato Bellotti , Romana Boiger  and Andreas Adelman ^{*} 

Paul Scherrer Institut, 5232 Villigen, Switzerland; renato.bellotti@psi.ch (R.B.); boiger.romana@gmail.com (R.B.)

^{*} Correspondence: andreas.adelmann@psi.ch

Abstract: Particle accelerators are enabling tools for scientific exploration and discovery in various disciplines. However, finding optimised operation points for these complex machines is a challenging task due to the large number of parameters involved and the underlying non-linear dynamics. Here, we introduce two families of data-driven surrogate models, based on deep and invertible neural networks, that can replace the expensive physics computer models. These models are employed in multi-objective optimisations to find Pareto optimal operation points for two fundamentally different types of particle accelerators. Our approach reduces the time-to-solution for a multi-objective accelerator optimisation up to a factor of 640 and the computational cost up to 98%. The framework established here should pave the way for future online and real-time multi-objective optimisation of particle accelerators.

Keywords: surrogate model construction; deep neural network; inverse neural network; charged particle accelerator; cyclotron; linear accelerator



Citation: Bellotti, R.; Boiger, R.; Adelman, A. Fast, Efficient and Flexible Particle Accelerator Optimisation Using Densely Connected and Invertible Neural Networks. *Information* **2021**, *12*, 351. <https://doi.org/10.3390/info12090351>

Academic Editor: Gianluca Valentino

Received: 16 July 2021

Accepted: 26 August 2021

Published: 28 August 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Advances in accelerator science and technology have enabled discoveries in particle physics and other fields—from chemistry and biology to medical applications—for more than a century [1], with no end in sight [2]. Particle accelerators consist of a multitude of building blocks, and the relationship between changes in machine settings, known as design variables, and the corresponding particle-beam response is often non-linear. For this reason, the development and operation of a particle accelerator rely heavily on computational models. These models use first principles of physics to state the equations of motion, while numerical algorithms are then employed to solve them. These models provide valuable insight, but their high computational cost is prohibitive for many applications. For example, a single simulation of the Argonne Wakefield Accelerator (AWA) (Figure 1) takes approximately ten minutes with the high-fidelity physics model Object-Oriented Parallel Accelerator Library (OPAL) [3], which renders such models unsuitable for real-time usage. In order to achieve optimal operation points of these machines, genetic algorithms (GAs) are typically the method of choice for solving multi-objective optimisation problems [4–8]. However, an optimisation sometimes requires thousands of model evaluations. For this reason, the time-to-solution can easily be in the order of days. Recent work [9] addressed this issue by training a data-driven surrogate model to approximate the computer model at a single position of interest. Evaluating such a surrogate model takes less than a second and results in predictions that are very close to the high-fidelity physics model.

Related to these works, Scheinker et al. developed and demonstrated an online multi-timescale multi-objective optimisation algorithm that performs real-time feedback on particle accelerators [10]. In contrast to our contribution, they demonstrate a general control theoretical motivated feedback algorithm, capable of performing online multi-objective optimisation. Bayesian optimisation in combination with particle-in-cell simulations [11] is used to tune a plasma accelerator autonomously to the beam energy spread to the sub-percent at a given energy and intensity.

In this work, we introduce two new surrogate models, a forward model and an invertible model. Both are designed to model the machine at any longitudinal position, overcoming a main limitation of existing approaches. Moreover, the proposed methods for building these models are general enough to be applicable to any kind of particle accelerators. We demonstrate the generality of the ansatz by considering a linear accelerator and a quasi-circular machine, i.e., a cyclotron. Furthermore, the fast surrogate models enable optimisations on time scales suitable for online and real-time multi-objective optimisations of particle accelerators.




Figure 1. Overview of the use cases. (a) Artistic representation of the IsoDAR experiment at KamLAND (Kamioka Observatory, Japan), from left to right the cyclotron, target and detector are depicted. (b) A schematic of the IsoDAR cyclotron with the relevant parameters. (c) The AWA facility at Argonne National Laboratory (US), pictured from the gun downstream (courtesy of J. Power). (d) Schematic of the relevant parts of the AWA machine.

The forward model resembles existing models described, e.g., in [9]. While previous work focused on approximating OPAL only at one position (usually at the end) of the accelerator (as in [9]), our new model provides an approximation to OPAL at a multitude of positions along the accelerator, allowing us to estimate beam properties along the entire accelerator depicted in Figure 1d. To achieve this goal, our model takes both the design variables and the position along the accelerator as input and predicts the beam properties at that position. Consequently, our model is a complete replacement for OPAL and does not limit our options of objectives and constraints for beam optimisations. Existing models are not capable of optimising the beam properties at multiple positions at the same time, but our forward model enables such sophisticated optimisations, which we demonstrate empirically.

The invertible model is capable of solving both the forward and the inverse problem by providing two kinds of predictions, called the forward and the inverse prediction. The

forward prediction approximates OPAL, i.e., it calculates the beam properties corresponding to a given machine configuration. The inverse prediction takes a user-imposed target beam and returns a machine configuration that realises an approximation to such a beam. In order to incorporate expert intuition and knowledge in the optimisation, the result of the inverse prediction initialises a GA-based multi-objective optimisation, which then requires fewer generations to converge than when the usual random initialisation is used.

We demonstrate the capabilities of the two surrogate models with the two real-world examples depicted in Figure 1. The first example is the AWA, which is a linear accelerator. The second one is the cyclotron used for the proposed Isotope Decay-at-Rest experiment, a proposed very-high-intensity electron-antineutrino source (from here on called the IsoDAR machine). We developed a forward model and an invertible model for both of these fundamentally different accelerators. Furthermore, we used them for a multi-objective accelerator optimisation to find machine configurations for desired beam properties.

1.1. Physics Models and Datasets

For both accelerators we used the same underlying physics model, based on the OPAL accelerator simulation framework. The dynamics of an accelerator are given by the following function:

$$\mathbf{f} : \mathbb{R}^m \times \mathbb{R} \rightarrow \mathbb{R}^n$$

$$\mathbf{f}(\mathbf{x}, s) = \mathbf{y}(s),$$

where the quantity $\mathbf{x}$ represents a vector of machine settings (design variables), the scalar $s \in \mathbb{R}$ denotes the position along the accelerator and the vector $\mathbf{y}(s) \in \mathbb{R}^n$ denotes the beam properties (quantities of interest). The function $\mathbf{f}$ represents the OPAL high-fidelity physics model.

Our surrogate models learn the relationship between design variables and quantities of interest based on examples, which means that we need to incorporate the physics of an accelerator into a dataset. We achieve this by sampling the space of design variables randomly and evaluating the sampled configurations of interest at multiple positions along the machine. One sample in the dataset consists of the design variables, a position along the machine and the corresponding quantities of interest.

1.2. Specifics of the Argonne Wakefield Accelerator Model

The AWA accelerator [12] consists mainly of an electron gun, several radio-frequency cavities for accelerating the electrons and magnets for keeping the beam focused. The electron gun generates bunches that form a train of Gaussians in the longitudinal direction, separated by a peak-to-peak distance λ . Other electron-gun variables are the gun phase ϕ , the charge Q , the laser spot size (SIGXY) and the current in the buck focusing and matching solenoids (BFS and MS). For details, we refer to Figure 1d. In total, we have nine design variables, denoted $\mathbf{x} \in \mathbb{R}^9$, that define the various operation points of the AWA. The ranges of $\mathbf{x}$ are determined by the physical limitations of the accelerator and listed in Table A3. The quantities of interest (QoIs), $\mathbf{y} \in \mathbb{R}^8$ are as follows: the transversal root-mean-square beam sizes σ_x, σ_y , the transversal normalised emittances ϵ_x, ϵ_y , the mean bunch energy E , the energy spread of the beam ΔE and the correlations between transversal position and momentum $\text{Corr}(x, p_x), \text{Corr}(y, p_y)$. A summary and more details about the used QoIs are given in Table A2.

We built a labeled dataset of the AWA by using OPAL and the latin hypercube procedure [13] in order to sample the search space of $\mathbf{x}$, which results in a training/validation set of 18,065 points and a test set of 913 points.

1.3. Specifics of the IsoDAR Cyclotron Model

The simulations are based on the latest IsoDAR [14–16] cyclotron design, with the nominal beam intensity of 5 mA H_2^+ beam (equivalent to 10 mA of protons). Ongoing

modeling efforts of IsoDAR consider the radial momenta p_{r0} , the injection radius r_0 , the phase ϕ_{rf} of the radio frequency of the acceleration cavities and the root mean square beam sizes $\sigma_x, \sigma_y, \sigma_z$ at injection. The physical ranges of the machine settings are given in Table A4. The surrogate models predict the transversal and longitudinal root mean square beam sizes $\sigma_x, \sigma_y, \sigma_z$, the projected emittances $\epsilon_x, \epsilon_y, \epsilon_z$, the beam halo parameters h_x, h_y, h_z , the energy E , the energy spread of the beam ΔE and the particle loss N_l (see Table A2). For this example, we have six machine settings $\mathbf{x} \in \mathbb{R}^6$ and twelve quantities of interest $\mathbf{y} \in \mathbb{R}^{12}$. All the samples are taken at turn 95, in the vicinity of the extraction channel.

The labeled dataset, obtained with OPAL, consists of 5000 random machine configurations. Among these, 1000 samples are randomly selected and used as the test set, 800 for the validation set and the remaining 3200 samples form the training set.

2. Results

2.1. Forward and Invertible Models

We assess the quality of the surrogate models by evaluating the adjusted coefficient of determination ($\bar{R}^2$) and the relative prediction error at 95% confidence on the test set; i.e., data that have not been used during the development of the models. For the AWA model, the $\bar{R}^2$ values for the forward model are very close to one at nearly all positions along the machine, except for the solenoids and cavities. As a consequence, the variance of the dataset is explained well, as can be observed in Figure 2a. Due to the presence of angular momenta inside the solenoid region, the variance of the emittance is not captured well. In OPAL this effect is not suppressed, however, Bush's theorem assures that, at the exit, the effect is reversed and does not contribute to our quantities of interest. In consequence, only the variance of the emittance is not captured well inside the solenoids and cavities, respectively. However, the relative prediction error for this quantity is still less than 25% for 95% confidence, as depicted in Figure 2c. The relative errors for the other quantities are even lower than 15% and 10% for σ_x (Figure 2b) and ΔE (Figure 2d), respectively.




Figure 2. Prediction metrics for the AWA forward surrogate model. The plots show the adjusted coefficient of determination (a) and the relative prediction error (b–d) as functions of the longitudinal position. All values are calculated on the test set, i.e., on samples that have not been considered when choosing the parameters. We omit the curves for the QOIs in the y direction for the sake of readability because they are close to the ones in x direction.

For the forward surrogate model of IsoDAR, the variance of the dataset is captured very well with a minimum $\bar{R}^2$ value of 0.95 among all predicted values, except for the

quantities in z-direction. Here, it ranges between 0.82 and 0.86. The relative prediction error of the forward model is at most 5.6%.

Since we are mainly interested in the beam parameters, we evaluated the invertible model by performing first the inversion, computing machine settings from beam parameters, with the invertible model and then a forward prediction with the forward model. For the AWA model the error at 95% confidence is almost 40% for the beam size, emittance and energy spread. Nevertheless, the sampling error is only 20% for three out of four target beams in the test set. The energy, on the other hand, is obtained almost with perfect accuracy.

For the IsoDAR model, the relative errors for 95% confidence range from 0.08% to 11%. The corresponding $\bar{R}^2$ values lie between 0.79 and 0.96 except for the longitudinal quantities, for which we obtained values between 0.65 and 0.77. The detailed values for the IsoDAR model can be found in Table A1, Figures A1 and A2.

Further details about the datasets and the model parameters can be found in Appendices B and C.

2.2. Multi-Objective Optimisation

To find good operation points for the accelerators, we solved multi-objective optimisation problems [5,6,9,17]. The standard approach for that is visualized in Figure 3a: A physics based model, such as OPAL, together with a GA and random initialization is used. We investigate two other approaches depicted in Figure 3b. First, we use a GA together with the forward model, instead of OPAL, and a random initialisation. In the second approach, the GA and the forward model are still employed, but the invertible model additionally provides us with a good initial guess for the optimisation.

2.3. Performance Metrics for the Optimisations

By characterising an approximation to the Pareto front, the set of optimal solutions encompasses two main aspects. First, we strive to converge quickly, using as few iterations as possible, towards non-dominated quantities of interest. The range of the objectives in the non-dominated set is chosen to analyse the convergence behaviour. Thus, for each generation the ranges of all objectives individually are compared for the two optimisation approaches. Second, we demand that the non-dominated quantities of interest are diverse, i.e., well spread out over the entire approximated Pareto front. The convex hull volume V_{ch} (see Appendix A.1 for more details), the number of solutions of the non-dominated set, as well as the generational distance [18] and the hypervolume difference (i.e., the difference between the hypervolume of the optimal solution and the hypervolumes of the Pareto front of each generation, with a reference point chosen close to the Nadir point [19]) assess this quality aspect. Finally, plotting projections of the non-dominating set provides a holistic view to compare the effect of the initialisation strategies (random vs. invertible model initialisation) on both the convergence and the diversity of the non-dominated set. These plots show the evolution of the solutions of the optimisation problem across the generations directly.




Figure 3. Optimisations with OPAL (a) and the forward model (b). The gray parts correspond to current practice, the blue parts belong to the genetic algorithm (GA) NSGA-II and the orange parts are our novel contributions. The GA generates n_g generations. Schematics of the forward model (c) and invertible model (d) architectures. The grey arrows denote the prediction directions.

2.4. Multi-Objective Optimisation for AWA

The constrained multi-objective optimisation problem of the following:

$$\begin{aligned}
 &\min && \sigma_x(26 \text{ m}) \\
 &\min && \epsilon_x(26 \text{ m}) \\
 &\min && \Delta E(26 \text{ m}) \\
 &\text{s. t.} && \sigma_x(s) \geq 2 \text{ mm}, \quad s \in 0, 0.25, 0.5, \dots, 10 \text{ m} \\
 &&& \sigma_x(s) \leq 5 \text{ mm}, \quad s \in 0, 0.25, 0.5, \dots, 10 \text{ m} \\
 &&& |\text{Corr}(x, p_x)| \leq 0.1, \quad \text{at } 26 \text{ m},
 \end{aligned}$$

represents a generic multi-objective optimisation task, found in many of the past and future AWA experiments [12]. The optimisation of beam parameters was required at 26 m for a pilot experiment performed in 2020, regarding a new scheme for electron cooling. In addition, there were multiple constraints at various positions along the accelerator. We solved the optimisation problem twice, once initialising the GA population randomly and once initialising the GA population with the invertible model outcome, where we utilized target values for the beam parameters as listed in Table A6.

The objective values for different generations are depicted in Figure 4a. The optimisation initialised with the invertible model converged almost entirely after 10 generations. If random initialisation is used, more than 100 generations are needed to arrive at the same optimal configurations. The ranges of the beam properties in the non-dominated set confirm this observation (Figure 4b–d). Using the invertible model results in a bigger hypervolume even when prediction uncertainty is accounted for (Figure 4e). Furthermore, the optimisation initialised with the invertible model finds more solutions during the first 500 generations (Figure 4f). These two statements imply that the initialisation with the invertible model finds both more and more diverse non-dominated design variables.

The non-dominated feasible machine settings are evaluated with OPAL (dots in Figure 5a–c) in order to validate the predicted optimal objective values. All the OPAL evaluated points lie within the region of uncertainty for the optimal objective values at 90% confidence.

In Figure 5d, deeper insights into the constraint on the correlation parameter ($\text{Corr}(x, p_x)$) is given. Although the surrogate models found many configurations lying outside the allowed region, both initialisation strategies result in at least some feasible configurations. Importantly, no machine configuration in the training/validation set is feasible according to the optimisation problem, that means that no sample in these datasets fulfills the optimisation constraints. Nevertheless, some optimal points fulfill the correlation constraint.




Figure 4. Convergence of the AWA optimisation. The figures show two-dimensional projections of the objective space (a), i.e., the quantities of interest at 26 m. The shaded area denotes the prediction uncertainty at 90% confidence, and the dots are the predictions themselves. Panels (b–d) depict the ranges of objective values among the non-dominated feasible solutions, the volume of the convex hull around the non-dominated solutions plus/minus Monte Carlo uncertainty (e) and the number of non-dominated solutions (f). The colour blue indicates that the optimiser is initialised randomly, and orange indicates that the invertible model is used for the initialisation. All quantities of interest are calculated using the forward surrogate model.




Figure 5. OPAL validation of the non-dominated feasible machine settings among all 1000 generations (dots) as calculated using the forward surrogate model. The colour blue marks optimisation with random initialisation, whereas orange indicates that the invertible model provided the initialisation. The dots are the results of an OPAL run, and the shaded areas mark values predicted by the forward model plus/minus uncertainty at 90% confidence.

2.5. Multi-Objective Optimisation for IsoDAR

For the IsoDAR case, we solved a two-objective optimisation problem, similar to that introduced by Edelen et al. [9]. We aim to minimise simultaneously the projected emittance ϵ_x and the energy spread ΔE of the beam without any constraints.

$$\begin{aligned} \min \quad & \Delta E \\ \min \quad & \epsilon_x. \end{aligned}$$

As before, the optimisation problem is solved with a GA by using the forward surrogate model. First, a random initialisation is used and then an initialisation using the invertible model with the target beam parameters in Table A5 is performed.

The convergence of the objective values for the two optimisation approaches is depicted in Figure 6a. Both approaches result in the same non-dominated set. The optimisation with the invertible model initialisation converges faster and finds more optimal points already in the first few generations. The random initialisation approach needs approximately 40 generations until it is at the same level. Figure 6b shows that both initialisation schemes result in the end to nearly the same number of solutions. The difference of the objective values for the two approaches in the first generations of the optimisation can also be observed in Figure 6d, which depicts the ranges from minimum to maximum values of the non-dominated set for the two objectives.

The performance-metrics plots Figure 6c,g show the behaviour of the two approaches quantitatively. The generational distance for the invertible model initialisation starts at a much lower value. Therefore, the non-dominated sets found in the first generation are already closer to the final non-dominated set than the ones found with the random initialisation. The lower hypervolume difference and the higher convex hull volume for the invertible model initialisation indicate that the initial non-dominated set is more widely spread, but these advantages vanish after approximately 40 generations.

In order to validate the optimisation results, we randomly chose 100 points of the non-dominated sets of both approaches and simulated them with OPAL. We did not evaluate all non-dominated machine settings because of the high computational cost of OPAL for the IsoDAR example. As can be seen in Figure 6e,f, nearly all of the OPAL points lie within the 90-percent confidence region of the prediction with the test set, although only few points of the original OPAL dataset are in this area.

2.6. Computational Advantages

We compare solving an optimisation problem with OPAL to our approach of using a forward surrogate model instead of OPAL. There are two quantities to be considered for the speedup analysis. First, the time-to-solution t measured in hours and, second, the computational cost c measured in CPU hours. The detailed calculations for the improvement of both quantities can be found in Appendix D in the Appendix. Previous work [9] also calculated speedups, but these calculations do not include the computational cost and the entire cost of developing a surrogate model because they neglect the hyperparameter scan. Moreover, the existing speedups refer to a model that is only capable of predicting the quantities of interest at one position. Our calculations refer to our novel forward surrogate model, which is capable of predicting quantities of interest at a plethora of positions and do not only include the generation of the dataset but also the training, hyperparameter tuning and the cost of running the optimisations themselves. The invertible model reduces the number of generations. However, we calculated the speedup for a fixed number of generations; therefore, the effect of the invertible model is not included in the speedup calculations.

For the AWA (see Table A8), we calculated the relative improvement $t_{\text{OPAL}}/t_{\text{Surr}} \approx 3.39$ in terms of time-to-solution and $c_{\text{OPAL}}/c_{\text{Surr}} \approx 1.83$ in terms of computational cost. This approach is more than three times faster and requires approximately 45% less computational resources (see Equation (A1) in Appendix D.3). For IsoDAR the benefits are even

more pronounced: $t_{\text{OPAL}}/t_{\text{surr}} \approx 640.00$ and $c_{\text{OPAL}}/c_{\text{surr}} \approx 59.20$. Hence, the usage of a surrogate model saves more than 98% of the computational cost and time resources compared to the approach using OPAL.




Figure 6. Convergence of the IsoDAR optimisation. The items coloured blue belong to the random initialisation, and the orange ones indicate the usage of the invertible model for initialisation. (a) Two-dimensional projection of the objective space for different generations. The dots are the predictions, and the shaded areas depict the prediction uncertainty at 90% confidence as estimated on the test set. The black circles show the initial values for the biased optimisation. (b) Number of non-dominated solutions. (c) Convergence metrics: Generational Distance (GD) (dots) and Hypervolume Difference (HV diff) (stars) to compare the two optimisation approaches. (d) Range (minimum/maximum) of the non-dominated set in objective space. (e,f) Validation with OPAL. Both figures show the objective space with OPAL validated points in black (triangles), the original OPAL data points in grey (triangles), the 90 percentile (shaded area) of the optimal Pareto front (blue and orange dots) and 100 randomly chosen points of the optimal Pareto front for validation (black dots). (g) Estimated volume of the convex hull of the non-dominated set configurations in objective space using the Monte Carlo uncertainty. The solid lines describe the mean values, and the shaded areas mark the mean plus/minus one standard deviation.

3. Methods

3.1. Forward Model

The forward surrogate model $\tilde{f}$ (Figure 3c) is a fast-to-evaluate approximation to the expensive OPAL-based physics model f .

$$\begin{aligned}\tilde{\mathbf{f}}: \mathbb{R}^m \times \mathbb{R} &\rightarrow \mathbb{R}^n \\ \tilde{\mathbf{f}}(\mathbf{x}, s) &= \tilde{\mathbf{y}}(s) \\ \text{s. t. } \tilde{\mathbf{f}}(\mathbf{x}, s) &\approx \mathbf{f}(\mathbf{x}, s) \quad \forall \mathbf{x}, s.\end{aligned}$$

While existing work [9] is restricted to approximating the accelerator at a single position s , usually at the end, we built a model for the AWA that approximates the dynamics of the machine at all positions. Without loss of generality, the IsoDAR model predicts the beam properties only at the end of the machine, as ongoing design work requires. We follow the work by Edelen et al. [9] and use densely connected feedforward neural networks [20] as candidates for the forward model. We focus only on architectures where all hidden layers are of the same width in order to simplify the hyperparameter scan. For the AWA model, the hidden layers use the Rectified Linear Unit (ReLU) activation and the output layer uses the tanh activation function, whereas for the IsoDAR model tanh is used as an activation for the hidden layers, and the output layer activation is linear. We use the mean absolute error (MAE) as the loss function for the AWA model and the mean squared error (MSE) for the IsoDAR model.

We optimise the trainable parameters by using the Adam algorithm with default parameters. The non-trainable parameters are selected using grid search and can be found in Table A7, along with further information about the models. All the models (forward and invertible) are implemented by using the TensorFlow framework [21] in version 2.0. For the hyperparameter scan of the IsoDAR model the RAY TUNE library [22] is used in addition. The models are trained and evaluated on the Merlin6 cluster at the Paul Scherrer Institut, using 12 CPU cores for the AWA dataset and one CPU core for the IsoDAR case.

3.2. Invertible Model

The invertible model (see Figure 3d) is capable of performing two kinds of predictions: A forward prediction approximating OPAL and an inverse prediction aiming to solve the following inverse problem. Given a target beam, the invertible model predicts a vector of design variables such that the corresponding beam is a good approximation of the target beam.

We follow the ansatz by Ardizzone et al. [23] and refer to their work for the details of the architecture of the inverse model. Ardizzone et al. build a neural network consisting of invertible layers, so-called affine coupling blocks (see Figure 3d). Each of them contains two neural networks, which are sharing parameters. We decided that all hidden layers of the internal networks have the same number of neurons, and the internal networks of all affine coupling blocks have the same architecture. Each internal network consists of the same number of hidden layers of the same size. All hidden neurons of the internal networks use the ReLU activation function.

The solution of the inverse problem is not unique, hence, a mechanism to select one solution is needed. This is realised by mapping the machine settings not only to the design variables but also to a latent space $\mathcal{Z} \subset \mathbb{R}^{d_z}$ of dimension d_z . This space follows a known probability distribution (e.g., uniform distribution as used in this work), which allows sampling random points in it. Sampling a point in the latent space corresponds to selecting one solution of the inverse problem. For this reason, the inverse prediction is also called sampling.

For technical reasons, an invertible neural network requires that the input and output vectors have the same length. This is generally not fulfilled. If the input and/or output vectors are not big enough, vectors containing noise of small magnitude are added so that the total dimension of input and output vectors is d . We follow the advice of

Ardizzone et al. [23] and allow padding not only on the smaller vector but also on both in order to increase the network width and, therefore, the model capacity. The total dimension d is a hyperparameter. We denote the padding vectors $\mathbf{x}_{\text{pad}} \in \mathbb{R}^p, \mathbf{y}_{\text{pad}} \in \mathbb{R}^q$ and obtain $d = m + 1 + p = n + 1 + \dim(\mathcal{Z}) + q$.

The mathematical formulation of the forward prediction is as follows.

$$\begin{aligned}\hat{\mathbf{f}}: \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^p &\rightarrow \mathbb{R}^n \times \mathbb{R} \times \mathcal{Z} \times \mathbb{R}^q \\ \hat{\mathbf{f}}(\mathbf{x}, s, \mathbf{x}_{\text{pad}}) &= (\hat{\mathbf{y}}, s, \mathbf{z}, \mathbf{y}_{\text{pad}}) \\ \text{s. t. } \hat{\mathbf{y}} &\approx \mathbf{f}(\mathbf{x}, s).\end{aligned}$$

The inverse prediction is now written as follows.

$$\begin{aligned}\hat{\mathbf{f}}^{-1}: \mathbb{R}^n \times \mathbb{R} \times \mathcal{Z} \times \mathbb{R}^q &\rightarrow \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^p \\ \hat{\mathbf{f}}^{-1}(\mathbf{y}, s, \mathbf{z}, \mathbf{y}_{\text{pad}}) &= (\hat{\mathbf{x}}, s, \mathbf{x}_{\text{pad}}) \\ \text{s. t. } \mathbf{f}(\hat{\mathbf{x}}, s) &\approx \mathbf{y}.\end{aligned}$$

In order to improve the performance of the invertible model we use a best-of- n strategy. For each inverse prediction (at inference time only), n_{tries} design variable configurations are sampled, evaluated with the forward pass and the best configuration is chosen as the final prediction. Note that we do not need to rely on a forward model but only on the invertible model itself. We chose $n_{\text{tries}} = 32$ for both the IsoDAR model and the AWA, as we observed no significant improvement for bigger values. All prediction errors are calculated with this particular choice. The loss function consists of multiple parts:

$$\mathcal{L}_{\text{inv}} = w_x \mathcal{L}_x + w_y \mathcal{L}_y + w_z \mathcal{L}_z + w_r \mathcal{L}_r + w_{\text{artificial}} \mathcal{L}_{\text{artificial}}, \quad (1)$$

where the $\mathcal{L}_x$ loss ensures that the sampled machine setting vectors $\hat{\mathbf{x}}$ follow the same distribution as the machine settings in the dataset, the $\mathcal{L}_z$ loss ensures that the latent space vectors follow the desired distribution. Both loss functions are realised as mean field discrepancy [24]. The $\mathcal{L}_y$ loss is the mean squared error between $\hat{\mathbf{y}}$ and $\mathbf{y}_{\text{true}}$, and the $\mathcal{L}_{\text{artificial}}$ loss is the sum of the MSEs of both $\mathbf{x}_{\text{pad}}$ and $\mathbf{y}_{\text{pad}}$. The reconstruction loss $\mathcal{L}_r$ ensures that the inverse prediction is robust with respect to perturbations of small amplitude ϵ .

$$\mathcal{L}_r = \sum_{i=1}^N \left\| \hat{\mathbf{f}}^{-1}(\hat{\mathbf{f}}(\mathbf{x}_i, s_i) + \epsilon) - \mathbf{x}_i \right\|^2.$$

The inverse prediction aims to generate machine settings $\hat{\mathbf{x}} = \hat{\mathbf{f}}^{-1}(\mathbf{y}, s)_{1:m}$ such that the resulting beam $\hat{\mathbf{y}} = \mathbf{f}(\hat{\mathbf{x}}, s)_{1:n}$ comes close to desired target-beam properties. In order to estimate the prediction accuracy of the models, we reproduce vectors $\mathbf{y}$ from the test set. It is not reasonable to compare the sampled machine settings with the corresponding setting $\mathbf{x}$ in the test set because multiple machine settings might realise similar beams. Instead, we performed a forward prediction with the forward surrogate model. This allows us to estimate which beam properties the sampled configuration results in i.e. $\tilde{\mathbf{y}} = \tilde{\mathbf{f}}(\hat{\mathbf{x}}, s)$. That quantity is then used to describe the error of the inverse prediction.

4. Discussion and Outlook

We have introduced a novel flexible forward surrogate model that is capable of simulating high-fidelity physics models at a plethora of positions along the accelerator. We have shown that forward surrogate models reduce both the time (by a speedup factor of 3.39 for the AWA and 640 for IsoDAR) and the computational cost (by 45% and 98%) of beam optimisations. The benefits increase if several optimisations need to be performed because the time-consuming and computationally expensive part is the development of the model. Once the models are developed, the optimisations are almost free and can be performed in approximately 10 min on 12 cores. This allows experimenters to adapt

to new situations by changing constraints and objectives and to rerun the optimisation with very little computational resources. Furthermore, the surrogate models themselves are useful in control-room settings, as they are fast enough to be used in a graphical user interface. This provides operators the possibility to quickly try out different machine settings. Despite these benefits of the surrogate model approach, one needs to be careful when applying them to problems with narrow constraints. Neural networks can only accurately model the machine-setting regimes represented in the dataset on which they are trained. If the feasible region of an optimisation problem is narrow, the models will hardly be able to find configurations that fulfill all constraints. One solution to this issue is to sample the feasible region more densely instead of sampling the design space uniformly. This is difficult because the feasible region in design space is usually a priori unknown. We have also presented an invertible model, which makes it possible to simulate the forward direction as well as the inversion, thus predicting machine settings that result in desired beam parameters. The invertible model alone has the potential to support the design of accelerators, and it could also be used to implement a failure-prediction system or an adaptive control system. Another use-case of invertible surrogate models was demonstrated in this paper. Invertible surrogate models were shown to be able to bias the initialisation of a GA towards the optimal region, thereby incorporating prior knowledge and experience of experimenters. This reduces the number of generations needed for convergence of the GA when a forward surrogate model is used.

As the optimisation is almost free when a forward surrogate is employed, there is not much incentive to go through the labour of developing an invertible model. However, if OPAL is used to evaluate all individuals, the reduced number of generations might reduce the time and cost of the optimisation significantly. Using OPAL solves the problem of the under-represented feasible space because OPAL is capable of modelling the relevant physics of a particle accelerator. The approach of applying the biased initialisation to an OPAL optimisation might combine both the advantage of needing fewer generations with the ability to accurately represent the feasible space. The accompanying savings might justify the cost of developing the invertible model. Further research is needed to investigate this prediction. Finally, the choice of the target vector is important, and additional research is needed to investigate its impact.

This research is a proof of concept and shows that the forward and inverse models of particle accelerators can be approximated with neural networks. So far, no real machine/experiment data are included in this work. In particular, the influence of measurement errors on the behavior of the neural network models is of great interest and needs further investigation. A first step to assess that would be to train the neural network models with simulated data and to test them with real measurement data. This would be possible only for the AWA but not for IsoDAR since the accelerator has not been built yet.

The societal impact of accelerator science and technology will continue, with no end in sight [1,2,16,25]. With this research study, we add new computational tools and, hence, contribute to the quest of finding optimal accelerator designs and machine configurations, which in turn are likely to greatly reduce construction and operational costs and to improve physics performance.

Author Contributions: R.B. (Renato Bellotti) and A.A. conceived the idea. R.B. (Renato Bellotti) and R.B. (Romana Boiger) implemented the models and optimisations and performed the analysis for AWA and IsoDAR, respectively. A.A. supervised the project. All authors wrote and edited the manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The experimental datasets used for this work are available through <https://doi.org/10.16907/eb4d0d49-9f29-48a2-899c-ced99ff3d8a1> (accessed on 16 June 2021). The code for the models is available at the following: https://gitlab.psi.ch/bellotti_r/invertible_networks_for_optimisation_of_pc (accessed on 16 June 2021).

Acknowledgments: We acknowledge the help of John Power from AWA and Daniel Winklehner from the IsoDAR collaboration.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Predicted Quantities and Model Fidelity

Appendix A.1. Performance Metrics

We make use of the adjusted coefficient of determination for assessing the quality of the surrogate models, which is defined as follows:

$$\bar{R}^2(s) := 1 - \frac{\sum_{i=1 \& s_i=s}^{N_{\text{test}}} (\tilde{\mathbf{f}}(\mathbf{x}_i, s_i) - \mathbf{y}_i)^2}{\sum_{i=1 \& s_i=s}^{N_{\text{test}}} (\mathbf{y}_i - \bar{\mathbf{y}})^2} \cdot \frac{N_{\text{test}} - 1}{N_{\text{test}} - m - 1},$$

where N_{test} denotes the number of samples in the test set, m is the number of design variables $\mathbf{x}_i \in \mathbb{R}^m$ and we only include samples where $s_i == s$. This quantity can be interpreted as the fraction of variance that is explained by the model. A perfect prediction corresponds to $\bar{R}^2 = 1$.

The prediction uncertainty at confidence q is estimated by calculating the residuals of the prediction and take the absolute value over the test set samples. Then, we calculate the q percentile of these values. The uncertainty is calculated separately for each beam property i and at every position s :

$$e_q = (q \text{ percentile at } s)_i |\tilde{y}_i - y_i|,$$

where the samples correspond to position s .

In order to measure the performance of the optimisation, we used among others the convex hull volume. To compute that, the residuals over the test set are approximated with a Gaussian distribution. This allows us to sample new points around the predicted values and to perform a Monte-Carlo estimate of the convex hull volume of the non-dominated set.

Appendix A.2. IsoDAR Model Fidelity

The performance of the IsoDAR forward model is shown in Figure A1 and Table A1.

The $\bar{R}^2$ value for the inverse prediction of the invertible IsoDAR model is shown in Table A1, and the error quantiles are depicted in Figure A2 of the Appendix.

	E	σ_x	σ_y	σ_z	h_x	h_y	h_z	ϵ_x	ϵ_y	ϵ_z	ΔE	N_I
50%	0	0.67	0.66	1.1	1.9	1.6	1.9	0.95	0.94	1.6	0.87	0.58
75%	0.01	1.2	1.1	1.9	3.2	2.8	3.3	1.7	1.8	2.6	1.5	0.99
90%	0.01	1.7	1.6	2.7	4.7	4	4.7	2.5	2.7	4	2.2	1.4
95%	0.02	2.1	2	3.2	5.5	5	5.6	3.3	3.4	5.2	2.6	1.7
99%	0.03	2.8	3.1	4.4	7.6	7.1	8	4.7	4.9	6.9	3.6	2.2

Figure A1. Percentiles of the relative prediction error for each quantity predicted by the forward IsoDAR model.

Table A1. $\bar{R}^2$ values of the IsoDAR models. The values of the row “inv FP” refer to the forward prediction of the invertible model, whereas the values in “inv IP” refer to its inverse prediction. The latter ones are calculated with the help of the forward model.

	E	σ_x	σ_y	σ_z	h_x	h_y	h_z	ϵ_x	ϵ_y	ϵ_z	ΔE	N_I
forward	1.0	0.97	0.96	0.85	0.9	0.95	0.86	0.99	0.99	0.82	0.98	1.0
inv FP	0.97	0.92	0.91	0.75	0.84	0.88	0.81	0.93	0.95	0.69	0.94	0.9
inv IP	0.96	0.94	0.88	0.77	0.83	0.82	0.77	0.91	0.91	0.65	0.91	0.79

	E	σ_x	σ_y	σ_z	h_x	h_y	h_z	ϵ_x	ϵ_y	ϵ_z	ΔE	N_l
50%	0.03	0.97	0.97	1.2	2.1	2.1	1.9	1.9	2.1	1.8	1.3	3.2
75%	0.05	1.6	1.7	2.1	3.8	4.1	3.6	3.9	3.7	3.1	2.6	5.8
90%	0.07	2.4	2.9	3.2	6.2	7.5	5.7	6.7	7	5.1	4.1	8.7
95%	0.08	3	4.1	4.2	7.7	10	7.3	8.4	9.4	6.9	5.6	11
99%	0.12	4.5	6.4	5.9	13	15	12	13	13	12	9.7	17

Figure A2. Prediction error for the inverse prediction of the IsoDAR model. This estimates how close the predicted machine settings approach the target beam. The rows represent the various confidence levels.

Appendix A.3. Model Summary

The quantities predicted by our models are summarised in Table A2.

Table A2. Quantities predicted by the various models. Checkmarks ✓ mean that the model predicts the quantity, while crosses x mean that the quantity is not predicted by the model.

	AWA		IsoDAR	
	Forward	Invertible	Forward	Invertible
E (MeV)	✓	✓	✓	✓
ΔE (MeV)	✓	✓	✓	✓
σ_x, σ_y (m)	✓	✓	✓	✓
ϵ_x, ϵ_y (mm mrad)	✓	✓	✓	✓
$\text{Corr}(x, p_x), \text{Corr}(y, p_y)$	✓	x	x	x
σ_z (m)	x	x	✓	✓
ϵ_z (mm mrad)	x	x	✓	✓
N_l	x	x	✓	✓
h_x, h_y, h_z	x	x	✓	✓

Appendix B. Parameter Ranges Used in the Optimisation

The parameter ranges for the design variables for the AWA and the IsoDAR models are given in Tables A3 and A4, respectively.

Table A3. The design variables and their ranges for the AWA models.

Bound	IBF (A)	IM (A)	ϕ (°)	ILS ₁ (A)	ILS ₂ (A)	ILS ₃ (A)	Q (nC)	λ (ps)	SIGXY (mm)
Lower	450	100	−50	0	0	0	0.3	0.3	1.5
Upper	550	260	10	250	200	200	5	2	12.5

Table A4. The design variables and their ranges for the IsoDAR accelerator models.

Bound	p_{r0} ($\beta\gamma$)	r_0 (mm)	ϕ_{rf} (°)	σ_x (mm)	σ_y (mm)	σ_z (mm)
Lower	0.002254	115.9	283.0	0.95	2.85	4.75
Upper	0.002346	119.9	287.0	1.05	3.15	5.25

From the 5000 samples for the IsoDAR model, 1000 were generated with machine configurations in the above ranges (Table A4). To increase the number of samples and improve the quality of the model, some ranges were narrowed. In detail, 500 samples were drawn with $\phi_{rf} \in [283.0^\circ, 284.0^\circ]$ and another 3500 samples with additionally $r_0 \in [116.9 \text{ mm}, 119.9 \text{ mm}]$.

Table A5. Target vector for the biased IsoDAR optimisation.

E	σ_x	σ_y	σ_z	h_x	h_y
112.0 MeV	2.4 mm	2.1 mm	1.5 mm	4.5	3.6
h_z	ϵ_x	ϵ_y	ϵ_z	ΔE	N_l
3.4	4.2 mm mrad	4.5 mm mrad	2.1 mm mrad	146.2 keV	7090.0

Table A6. Target vector for the biased AWA optimisation.

E	σ_x	σ_y	ϵ_x	ϵ_y	ΔE	s
47.5 MeV	2 mm	2 mm	3 mm mrad	3 mm mrad	60 keV	13.7 m

Appendix C. Model Parameters

All parameters and properties concerning the neural network models are listed in Table A7.

Explanation of the network architectures: A feedforward network with n_l hidden layers of width n_w is denoted as an $n_l \times n_w$ network. An invertible model consisting of n_b blocks, each containing internal networks of depth n_d and width n_w , is denoted as $n_b \times n_d \times n_w$.

The loss functions were described in the main text and are given here for completeness reasons: $\text{MAE} = \frac{1}{N} \sum_{i=1}^N \|\tilde{\mathbf{f}}(\mathbf{x}_i, s_i) - \mathbf{y}_i\|$, $\text{MSE} = \frac{1}{N} \sum_{i=1}^N \|\tilde{\mathbf{f}}(\mathbf{x}_i, s_i) - \mathbf{y}_i\|^2$, where N denotes the number of samples over which the loss is calculated.

Table A7. Parameters and properties of our surrogate models. EOM means end of the machine.

	AWA		IsoDAR	
	Forward	Invertible	Forward	Invertible
Modelled region	$s \in [0, 26]$ m (entire machine)	$s \in [0, 13.7]$ m	EOM	EOM
Preprocessing $\mathbf{x}$	Scale to $[-1, 1]$	QuantileScaler Scale to $[-1, 1]$	Scale to $[0, 1]$	Scale to $[0, 1]$
Preprocessing $\mathbf{y}$	Shift to be positive Apply $\log()$ Scale to $[-1, 1]$	Clip $\epsilon_x \in [0, 200]$ mm mrad Scale to $[-1, 1]$	Scale to $[0, 1]$	Scale to $[0, 1]$
Dimension of the latent space	-	1	-	1
Nominal Dimension	-	12	-	14
Distribution of the latent space	-	Unif(-1, 1) $\mathcal{L}_{\text{inv}}$ (Equation (1)) with weights: $w_x = 400$ $w_y = 400$ $w_z = 400$ $w_r = 3$ $w_{\text{artificial}} = 1$	-	Unif(-1, 1) $\mathcal{L}_{\text{inv}}$ (Equation (1)) with weights: $w_x = 400$ $w_y = 400$ $w_z = 400$ $w_r = 10$ $w_{\text{artificial}} = 1$
Loss function	MAE		MSE	
Training algorithm	Adam	Adam	Adam	Adam
Learning rate	10^{-4}	10^{-4}	10^{-4}	10^{-3}
Batch size	256	256	256	8
Number of epochs	56	15	5000	30
Architecture	7×500	$8 \times 3 \times 100$	6×80	$5 \times 2 \times 60$
Activation of hidden neurons	ReLU	ReLU	tanh	ReLU
Number of trainable parameters	1,512,618	688,192	33,932	91,340
CPU cores for training	12	12	1	1
Time for training	49 h	31 h	1 h	1 h

Appendix D. Computational Details

Appendix D.1. Hardware

All computations were performed on the Merlin6 cluster at PSI. Each cluster node contains two Intel® Xeon® Gold 6152 Processors, each of which provides 22 cores (2 threads per physical core). Each node provides 384 GB of memory.

Appendix D.2. Implementation

The forward model and invertible model are implemented in TensorFlow 2.0 [21]. They are trained on 12 CPU cores for the AWA dataset and one core for the IsoDAR model. Training and evaluation were executed on the Merlin6 cluster at the Paul Scherrer Institute. The hyperparameters of IsoDAR are found with a grid search using [22].

We solve both the AWA and IsoDAR optimisation problems using the NSGA-II algorithm [26] implemented by the Python library pymoo [27], with default parameters. The general algorithm is shown in Algorithm A1. The NSGA2() function is given by the quantities of interest for the current generation, calculates the values of the objectives and the constraints and suggests the next generation of machine settings based on the results. When we refer to the optimal solutions of generation g , we refer to the non-dominated feasible solutions that are found in the generations up to and including generation g . Notice that only the forward surrogate model is used to evaluate the objectives and constraints; OPAL is used exclusively to train the models and to validate the final optimal configurations.

We solved both optimisation problems in two ways: First, we initialised the first generation randomly by uniformly sampling machine settings from their ranges (see Tables A3 and A4). Second, the invertible model is used to bias the initialisation towards the optimal region. To achieve this, we provided a target vector of beam properties $\mathbf{y}_t$ at position s_t and asked the invertible model to sample the individuals in order to achieve this beam, see Algorithm A2. In other words, we guessed what a good beam would look like at one position, and let the network calculate the corresponding machine settings. This is what allows operators and experimenters to incorporate their experience and intuition into the optimisation.

Appendix D.3. Speedup Calculation

Let t_o be the time needed for a single OPAL simulation, running on r_o cores. Let t_t be the time to train a surrogate model on r_t CPU cores, and t_p the time needed to evaluate one individual of the optimisation (the p stands for prediction). Assume that the prediction takes place on the same number of cores as the training. Let n be the number of unique machine configurations in the datasets. Since the development of a surrogate model involves a training, validation and test set, all of them are included in this number. Let n_h be the number of hyperparameter configurations to be tried for the model. A single optimisation requires running n_g generations, each consisting of n_i individuals. A summary of the parameters along with their values for both models is depicted in Table A8. The runtime of OPAL depends on the machine settings. The quantity is obtained by measuring the runtime for generating the training/validation set and then dividing by the number of unique machine settings in this set.

Algorithm A1: Optimisation using surrogate models.

Input: Trained forward surrogate $\tilde{\mathbf{f}}$.
Input: Number of generations N_{gen} .
Input: Number of individuals in each generation N_{ind} .
Input: Positions where to evaluate the objectives $\mathcal{S}_{\text{obj}}$ and constraints $\mathcal{S}_{\text{con}}$.
Input: NSGA-II function to advance by one generation $\text{NSGA2}()$.
Input: Function that determines all non-dominated solutions from a set $\text{non_dominated}()$.
Result: Optimal set of machine settings X^* .
Result: Optimal objectives as predicted by the forward surrogate $\tilde{Y}^*$.
Result: Optimal objectives as calculated by OPAL Y^* .

```

begin
   $X_0 \leftarrow \text{initialise}(N_{\text{ind}})$ 
   $X \leftarrow \{\}$ 
   $\tilde{Y} \leftarrow \{\}$ 
  for  $i \in \{1, \dots, N_{\text{gen}}\}$  do
    Evaluate the current generation:
     $\tilde{Y}_i \leftarrow \{\}$ 
    for  $\mathbf{x} \in X_{i-1}$  do
      for  $s \in \mathcal{S}_{\text{obj}} \cup \mathcal{S}_{\text{con}}$  do
         $\tilde{\mathbf{y}}_{i,s} \leftarrow \tilde{\mathbf{f}}(\mathbf{x}, s)$ 
         $\tilde{Y}_i \leftarrow \tilde{Y}_i \cup \{\tilde{\mathbf{y}}_{i,s}\}$ 
      end
    end
     $\tilde{Y} \leftarrow \tilde{Y} \cup \tilde{Y}_i$ 
    Build the new generation:
     $X_i \leftarrow \text{NSGA2}(\tilde{Y}_i)$ 
     $X \leftarrow X \cup X_i$ 
  end
  Select the non-dominated individuals:
   $X^* \leftarrow \text{non\_dominated}(X)$ 
   $\tilde{Y}^* \leftarrow \{\tilde{\mathbf{f}}(\mathbf{x}, s) \mid \mathbf{x} \in X^*, s \in \mathcal{S}_{\text{obj}}\}$ 
  Calculate the optimal beam parameters with OPAL:
   $Y^* \leftarrow \{\mathbf{f}(\mathbf{x}, s) \mid \mathbf{x} \in X^*, s \in \mathcal{S}_{\text{obj}}\}$ 
end
  
```

Algorithm A2: Biased initialisation using the invertible surrogate model.

Input: Trained invertible surrogate $\hat{\mathbf{f}}$.
Input: Number of individuals per generation N_{ind} .
Input: Target vector $\mathbf{y}_t$ at position s_t .
Result: First generation X_0 .

```

begin
   $X_0 = \{\}$ 
  for  $i \in \{1, \dots, N_{\text{ind}}\}$  do
     $\mathbf{z} \leftarrow \text{Unif}(-1, 1)$ 
     $\mathbf{y}_{\text{pad}} \leftarrow 5 \cdot 10^{-2} \cdot \text{Normal}(0, 1)$ 
     $\mathbf{x} \leftarrow \hat{\mathbf{f}}^{-1}(\mathbf{y}_t, s_t, \mathbf{z}, \mathbf{y}_{\text{pad}})$ 
     $X_0 \leftarrow X_0 \cup \{\mathbf{x}\}$ 
  end
end
  
```

For all calculations, assume that we have infinite computational resources at our disposal. This means that all calculations can always exhaust the entire theoretically available parallelism. In practice, run times are usually higher than the ones calculated using this assumption. The only step where the surrogate model is limited by computational resources is the generation of the dataset. This step is embarrassingly parallel, so the surrogate model can adapt perfectly to limited computational resources. If only OPAL is used, however, the optimisation is also affected by limited resources. In the case where only one individual of a generation cannot be evaluated in parallel to the others, the time needed for the optimisation already doubles. This is because the next generation of the optimisation algorithm cannot start before the parent generation is fully evaluated (assuming that the optimisation algorithm features no parallelism across generations, which is the case for the pymoo implementation). Therefore, our assumption of infinite parallelism favours the OPAL-only approach. For this reason, our calculations result in a lower bound for the speedup.

Table A8. Parameters related to the speedup calculation. For the AWA, the prediction time refers to predicting the beam properties with a spacing of 5 cm, i.e., at 520 positions. For the IsoDAR accelerator, we modelled only 1 position. The times are measured by measuring the prediction time for predicting N_m machines and then dividing by N_m . For the AWA case, $N_m = 64$; for IsoDAR $N_m = 2^{20}$.

Quantity	Symbol	AWA	IsoDAR
Time for one OPAL evaluation	t_o	10 min	1.8 h
Time to train one forward model	t_t	49 h	1 h
Time to predict one machine with the surrogate	t_p	21 ms	26 μ s
Number of unique machine settings in the dataset	n	21.000	5.000
Number of hyperparameters to try	n_h	100	120
Number of CPU cores per OPAL evaluation	r_o	4	1
Number of CPU cores to train and evaluate the surrogate	r_t	12	1
Number of generations	n_g	1.000	1.000
Number of individuals per generation	n_i	200	300

Now we can calculate the time-to-solution for an optimisation using the surrogate model. This task involves the creation of the dataset and the model training (including a hyperparameter scan). The former is embarrassingly parallel, and we have unlimited computational resources; thus, all design variable configurations are evaluated at once. The same argument applies for the hyperparameter scan. The time-to-solution for an optimisation with the surrogate model is given by

$$t_{\text{surr}} = t_o + t_t + n_g \cdot t_p.$$

The associated computational cost is calculated as follows

$$c_{\text{surr}} = n \cdot t_o \cdot r_o + n_h \cdot t_t \cdot r_t + n_g \cdot n_i \cdot t_p \cdot r_t.$$

The time-to-solution for an OPAL optimisation is calculated as

$$t_{\text{OPAL}} = n_g \cdot t_o.$$

The corresponding computation cost is calculated by the following

$$c_{\text{OPAL}} = n_g \cdot n_i \cdot t_o \cdot r_o.$$

Now, we can determine the speedup in terms of execution time:

$$\begin{aligned}\frac{t_{\text{OPAL}}}{t_{\text{surr}}} &= \frac{n_g \cdot t_o}{t_o + t_t + n_g \cdot t_p} \\ &= \frac{n_g}{1 + \frac{t_t}{t_o} + n_g \cdot \frac{t_p}{t_o}} \\ &\approx \frac{n_g}{1 + \frac{t_t}{t_o}},\end{aligned}$$

where we have used $n_g \cdot \frac{t_p}{t_o} \ll 1$ for the approximation. This formula can be interpreted as follows: the greater the runtime of OPAL compared to the training time and the more generations we want to calculate for the optimisation, the higher the speedup will be.

The improvement in computational cost can be calculated by

$$\begin{aligned}C &= \frac{c_{\text{OPAL}}}{c_{\text{surr}}} \\ &= \frac{n_g \cdot n_i \cdot t_o \cdot r_o}{n \cdot t_o \cdot r_o + n_h \cdot t_t \cdot r_t + n_g \cdot n_i \cdot t_p \cdot r_t}.\end{aligned}$$

The relative reduction in computational resources can then be calculated by

$$\frac{c_{\text{surr}} - c_{\text{OPAL}}}{c_{\text{OPAL}}} = \frac{c_{\text{surr}}}{c_{\text{OPAL}}} - 1 = \frac{1}{C} - 1. \quad (\text{A1})$$

Solving n optimisation problems is equivalent to solving one problem using n times the number of generations. This means that solving more optimisation problems is equivalent to using more generations. Asymptotically, the improvement in terms of computational cost scales like

$$C \xrightarrow{n_g \rightarrow \infty} \frac{t_o}{t_p} \cdot \frac{r_o}{r_t}.$$

The relationship of the relative improvement and the number of optimisations can be observed in Figure A3.




Figure A3. Relative improvement in terms of time-to-solution and computational cost for the AWA (a) and the IsoDAR (b) optimisations. The curves are calculated analytically based on measurements of the prediction time for the surrogate model and OPAL.

References

- Shiltsev, V. Particle beams behind physics discoveries. *Phys. Today* **2020**, *73*, 32. [CrossRef]
- No final frontier. *Nat. Rev. Phys.* **2019**, *1*, 231. [CrossRef]
- Adelmann, A.; Calvo, P.; Frey, M.; Gsell, A.; Locans, U.; Metzger-Kraus, C.; Neveu, N.; Rogers, C.; Russell, S.; Sheehy, S.; et al. OPAL a Versatile Tool for Charged Particle Accelerator Simulations. *arXiv* **2019**, arXiv:1905.06654.
- Emery, L. Global Optimization of Damping Ring Designs using a Multi-objective Evolutionary Algorithm. In Proceedings of the Particle Accelerator Conference, Knoxville, TN, USA, 16–20 May 2005; p. 2962.
- Gulliford, C.; Bartnik, A.; Bazarov, I.; Maxson, J. Multiobjective optimization design of an rf gun based electron diffraction beam line. *Phys. Rev. Accel. Beams* **2017**, *20*, 033401. [CrossRef]

6. Neveu, N.; Spentzouris, L.; Adelmann, A.; Ineichen, Y.; Kolano, A.; Metzger-Kraus, C.; Bekas, C.; Curioni, A.; Arbenz, P. Parallel general purpose multiobjective optimization framework with application to electron beam dynamics. *Phys. Rev. Accel. Beams* **2019**, *22*, 054602. [\[CrossRef\]](#)
7. Frey, M.; Snuverink, J.; Baumgarten, C.; Adelmann, A. Matching of turn pattern measurements for cyclotrons using multiobjective optimization. *Phys. Rev. Accel. Beams* **2019**, *22*, 064602. [\[CrossRef\]](#)
8. Kranjčević, M.; Gorgi Zadeh, S.; Adelmann, A.; Arbenz, P.; van Rienen, U. Constrained multiobjective shape optimization of superconducting rf cavities considering robustness against geometric perturbations. *Phys. Rev. Accel. Beams* **2019**, *22*, 122001. [\[CrossRef\]](#)
9. Edelen, A.; Neveu, N.; Frey, M.; Huber, Y.; Mayes, C.; Adelmann, A. Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems. *Phys. Rev. Accel. Beams* **2020**, *23*, 044601. [\[CrossRef\]](#)
10. Scheinker, A.; Hirlaender, S.; Velotti, F.M.; Gessner, S.; Della Porta, G.Z.; Kain, V.; Goddard, B.; Ramjiawan, R. Online multi-objective particle accelerator optimization of the AWAKE electron beam line for simultaneous emittance and orbit control. *AIP Adv.* **2020**, *10*, 055320. [\[CrossRef\]](#)
11. Jalas, S.; Kirchen, M.; Messner, P.; Winkler, P.; Hübner, L.; Dirkwinkel, J.; Schnepf, M.; Lehe, R.; Maier, A.R. Bayesian Optimization of a Laser-Plasma Accelerator. *Phys. Rev. Lett.* **2021**, *126*, 104801. [\[CrossRef\]](#)
12. Power, J. Advanced Acceleration Concepts at the Argonne Wakefield Accelerator Facility. APS Division of Physics of Beams Annual Newsletter. 2018; p. 21. Available online: <https://www.anl.gov/awa> (accessed on 2 July 2021).
13. McKay, M.D.; Beckman, R.J.; Conover, W.J. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics* **1979**, *21*, 239. [\[CrossRef\]](#)
14. Bungau, A.; Adelmann, A.; Alonso, J.R.; Barletta, W.; Barlow, R.; Bartoszek, L.; Calabretta, L.; Calanna, A.; Campo, D.; Conrad, J.M.; et al. Proposal for an Electron Antineutrino Disappearance Search Using High-Rate ⁸Li Production and Decay. *Phys. Rev. Lett.* **2012**, *109*, 141802. [\[CrossRef\]](#) [\[PubMed\]](#)
15. Waites, L.H.; Alonso, J.R.; Conrad, J. IsoDAR: A cyclotron-based neutrino source with applications to medical isotope production. *AIP Conf. Proc.* **2019**, *2160*, 040001. [\[CrossRef\]](#)
16. Alonso, J.R.; Barlow, R.; Conrad, J.M.; Waites, L.H. Medical isotope production with the IsoDAR cyclotron. *Nat. Rev. Phys.* **2019**, *1*, 533–535. [\[CrossRef\]](#)
17. Bazarov, I.V.; Sinclair, C.K. Multivariate optimization of a high brightness dc gun photoinjector. *Phys. Rev. ST Accel. Beams* **2005**, *8*, 034202. [\[CrossRef\]](#)
18. Veldhuizen, V.; Allen, D. *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*; Technical Report, Evolutionary Computation; Air Force Institute of Technology: Dayton, OH, USA, 1999.
19. Fonseca, C.M.; Paquete, L.; Lopez-Ibanez, M. An Improved Dimension-Sweep Algorithm for the Hypervolume Indicator. In Proceedings of the 2006 IEEE International Conference on Evolutionary Computation, Vancouver, BC, Canada, 16–21 July 2006; pp. 1157–1163. [\[CrossRef\]](#)
20. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016. Available online: <http://www.deeplearningbook.org> (accessed on 16 June 2021).
21. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G.S.; Davis, A.; Dean, J.; Devin, M.; et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. 2015. Available online: [tensorflow.org](https://www.tensorflow.org) (accessed on 3 August 2021).
22. Liaw, R.; Liang, E.; Nishihara, R.; Moritz, P.; Gonzalez, J.E.; Stoica, I. Tune: A Research Platform for Distributed Model Selection and Training. *arXiv* **2018**, arXiv:1807.05118.
23. Ardizzone, L.; Kruse, J.; Wirkert, S.; Rahner, D.; Pellegrini, E.W.; Klessen, R.S.; Maier-Hein, L.; Rother, C.; Köthe, U. Analyzing Inverse Problems with Invertible Neural Networks. *arXiv* **2018**, arXiv:1808.04730v3.
24. Gretton, A.; Borgwardt, K.M.; Rasch, M.J.; Schölkopf, B.; Smola, A. A kernel two-sample test. *J. Mach. Learn. Res.* **2012**, *13*, 723–773.
25. Abs, M.; Adelmann, A.; Alonso, J.R.; Axani, S.; Barletta, W.A.; Barlow, R.; Bartoszek, L.; Bungau, A.; Calabretta, L.; Calanna, A.; et al. IsoDAR@ KamLAND: A Conceptual Design Report for the Technical Facility. *arXiv* **2015**, arXiv:1511.05130.
26. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **2002**, *6*, 182–197. [\[CrossRef\]](#)
27. Blank, J.; Deb, K. Pymoo: Multi-Objective Optimization in Python. *IEEE Access* **2020**, *8*, 89497–89509. [\[CrossRef\]](#)