

Machine learning applied to proton radiography of high-energy-density plasmas

Nicholas F. Y. Chen,¹ Muhammad Firmansyah Kasim,² Luke Ceurvorst,¹ Naren Ratan,¹ James Sadler,¹ Matthew C. Levy,¹ Raoul Trines,³ Robert Bingham,³ and Peter Norreys^{1,3}

¹*Clarendon Laboratory, Department of Physics, University of Oxford, Oxford OX1 3PU, United Kingdom*

²*John Adams Institute, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom*

³*STFC Rutherford Appleton Laboratory, Chilton, Didcot OX11 0QX, United Kingdom*

(Received 3 September 2016; published 17 April 2017)

Proton radiography is a technique extensively used to resolve magnetic field structures in high-energy-density plasmas, revealing a whole variety of interesting phenomena such as magnetic reconnection and collisionless shocks found in astrophysical systems. Existing methods of analyzing proton radiographs give mostly qualitative results or specific quantitative parameters, such as magnetic field strength, and recent work showed that the line-integrated transverse magnetic field can be reconstructed in specific regimes where many simplifying assumptions were needed. Using artificial neural networks, we demonstrate for the first time 3D reconstruction of magnetic fields in the nonlinear regime, an improvement over existing methods, which reconstruct only in 2D and in the linear regime. A proof of concept is presented here, with mean reconstruction errors of less than 5% even after introducing noise. We demonstrate that over the long term, this approach is more computationally efficient compared to other techniques. We also highlight the need for proton tomography because (i) certain field structures cannot be reconstructed from a single radiograph and (ii) errors can be further reduced when reconstruction is performed on radiographs generated by proton beams fired in different directions.

DOI: [10.1103/PhysRevE.95.043305](https://doi.org/10.1103/PhysRevE.95.043305)

I. INTRODUCTION

Magnetic fields generated in laser-matter interactions are of primary interest in high-energy-density physics [1]. For example, magnetic fields generated by the Weibel instability can explain the collisionless shocks that are found in young galaxies and other astrophysical systems [2,3]. In inertial confinement fusion, magnetic fields are used in one approach to reduce heat losses and thus improve performance of implosions [4], and in another approach (using cylindrical implosions) as a necessary criterion to reach ignition [5]. Also, magnetic reconnection is a commonly studied process, which converts some of the magnetic energy of a system into heat, and understanding the heating mechanism well could lead to better hohlraum design for inertial confinement fusion [6].

Proton radiography is an extensively used technique that characterizes electric and magnetic fields in plasmas over a wide range of field strengths [7]. A polyenergetic proton beam, with typical energies on the order of 10 MeV, is usually produced by high-intensity laser interaction with solid targets [8]. This beam then interacts with an object of interest (such as plasmas or shock-compressed matter) and gets deflected as a result of the Lorentz force or collisions with atoms [9]. The outgoing beam is captured on a radiochromic film (RCF) stack, which can resolve both spatial and energy profiles [10] (see Fig. 1 for an example of an experimental radiograph). In this paper we consider the case where the proton only interacts with the magnetic fields of high-energy-density plasmas via the Lorentz force, though our method also applies to other forms of scattering.

Various methods have been developed in analyzing proton radiographs. Using the principles of differential scattering and stopping, density profiles of dense matter can be retrieved from radiographs [11]. Via scaling laws, field strengths of electric and magnetic fields can be estimated [12–14]. Also, radiographs can be used to qualitatively understand electric and

magnetic field structures [15–17]. Furthermore, radiographs can be simulated numerically in order to identify features found in experimental radiographs [18,19].

It is only recently that techniques have been developed to reconstruct fields. The relations between the field structures and proton radiographs have been established by Kugland *et al.* [20] under certain simplifying assumptions, allowing one to obtain the line-integrated transverse magnetic field from a radiograph by solving a 2D Poisson equation. Graziani *et al.* [21] and Kasim *et al.* [22] provided extensions to this technique, under similar assumptions. As such, radiographs of systems that do not obey any of the assumptions in Refs. [20–22] can only be analyzed qualitatively.

Machine learning, a field of study which enables the performance of a computer (with respect to a certain task) to increase with its experience, has seen many applications in artificial intelligence problems such as image recognition, recommender systems, and speech-to-text [23]. Due to its ability to discover structures in high-dimensional data, artificial neural networks (one example of machine learning) has seen many applications in physics, such as analyzing particle accelerator data [24], reconstructing images in optical tomography [25], and retrieving 3D potentials in electron scattering [26]. The flexible nature of artificial neural networks and the prevalence of its usage in image recognition problems prompt us to posit its usage in imaging 3D magnetic field structures without a need for simplifying assumptions, addressing the gaps found in existing radiograph inversion techniques.

In this paper, we first review existing work on inverting proton radiographs. We then introduce key ideas of artificial neural networks and review their applications in physics. Next, we outline our new method of using artificial neural networks to reconstruct, for the first time, 3D magnetic fields in the nonlinear regime. We also suggest a simple way of retrieving field parameters such as characteristic lengths. Via simulations, we show a proof of concept for the above ideas and discuss how




FIG. 1. Sample of an experimental radiograph. The whiter regions indicate areas with higher proton intensity.

noise and selection of training data affect our results. Using an example, we highlight the need for proton tomography. Finally, we compare the artificial neural network technique with the existing methods of radiograph inversion and suggest a variety of extensions to our research.

II. THEORY

A. Existing methods of retrieving magnetic fields from radiographs

In this subsection, we will outline the foundational work on proton radiograph inversion by Kugland *et al.*, move on to discuss Graziani *et al.* and Kasim *et al.*'s extensions, and conclude with the gaps in these methods.

First, we go through Kugland *et al.*'s [20] definitions: The coordinates are defined such that the object is placed at $z = 0$ (object plane), and (x, y) refers to the coordinates on the image plane (see Fig. 2). At the object plane, the proton's coordinates are denoted as (x_0, y_0) . The distance between the proton source and the object is l while the distance between the object to the image plane (radiochromic film stack) is L . a , the characteristic

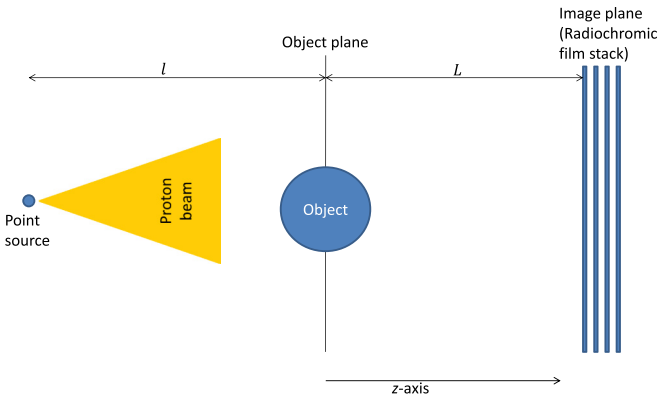


FIG. 2. Diagram of a typical proton radiography setup. A point source of distance l away from the object emits a beam of protons moving generally in the z direction. L is the distance between the object and image plane.

length of the object, is assumed to be much smaller than l (paraxial limit) and $L \gg l$ for high magnification.

In order to get a tractable result, Kugland *et al.* have made some simplifying assumptions. We start off with those relating to the proton source: (i) The source can be treated as a point source. Else, the radiograph will be blurred and the resolution of field structures will be affected. (ii) The protons deviate from their straight-line trajectories solely due to the Lorentz force interaction with the object, and we can ignore space-charge effects because the beam is charge-neutral as a result of comoving electrons [27]. (iii) The angular width of the beam is much greater than a/l so that intensity variations in the image plane are due to proton interactions with the object, and not the angular distribution of the proton beam.

Consider the dimensionless parameter

$$\mu \equiv \frac{l\beta}{a}, \quad (1)$$

where a is a characteristic length of the electromagnetic field, and β is a characteristic deflection angle. One core assumption in Kugland *et al.* is that $\mu \ll 1$ (hence known as the linear regime), where the spatial variation of the intensity on the screen is small. This is in contrast to the nonlinear regime (μ on the order of 1 or more) where the intensity variations are large, leading to nonlinear features. One example of nonlinear features is caustics, which occurs when the Jacobian determinant

$$\left| \frac{\partial(x, y)}{\partial(x_0, y_0)} \right| = 0, \quad (2)$$

resulting in features of high intensity (usually multiples of the background intensity).

Furthermore, assuming that the velocity of the proton v is approximately constant while the proton is within the object (trajectories are not perturbed within the plasma so $dt = dz_0/v$), the only relevant component of the magnetic potential is the one in the z direction, A_z . Defining the line-integrated potential as

$$\Phi(x_0, y_0) = \int_{-\infty}^{\infty} A_z(x_0, y_0, z_0) dz_0, \quad (3)$$

then with all the assumptions listed above, Kugland *et al.*'s formula for radiograph inversion reads

$$\nabla_{\perp}^2 \Phi(x_0, y_0) = \frac{\sqrt{2m_p K}}{el} \left(1 - \frac{I L^2}{I_0 l^2} \right), \quad (4)$$

where ∇_{\perp} is the gradient with respect to the transverse coordinates (x_0, y_0) , m_p is the mass of the proton, K is the (nonrelativistic) kinetic energy of the proton, e is the charge of an electron, I is the proton intensity distribution at the image plane, and I_0 is the proton intensity distribution in the object plane. As such, given the intensity profile at the object plane $I_0(x_0, y_0)$ and radiograph intensity profile $I(x, y)$ (which can be transformed to $I(x_0, y_0)$ via the mapping $x = \frac{l}{L}x_0, y = \frac{l}{L}y_0$) in the regime $\mu \ll 1$, one can solve a 2D Poisson equation to get the line-integrated potential $\Phi(x_0, y_0)$, thereby allowing one to reconstruct the line-integrated transverse magnetic field by taking the curl.

Using a series of perturbations, and assuming the linear regime $\mu \ll 1$, Graziani *et al.* [21] proposed a correction term

in the right-hand side of Eq. (20) in Kugland *et al.* (Eq. (4) in this paper), which leads to a second-order nonlinear partial differential equation. The authors then conducted a simulation of their proposed equation and found that their method reconstructed the line-integrated magnetic field accurately at locations near the peak field strength but was inaccurate at locations where the field strengths are at least 3 orders of magnitude less than the peak field strength. Also, Graziani *et al.* briefly sketched a method to retrieve the line-integrated magnetic field in the nonlinear regime, assuming that the direction of the proton trajectory within the object is nearly constant.

Another method, based on computational geometry, was implemented by Kasim *et al.* [22]. This method works well in the beginning of the caustic regime (early part of the nonlinear regime where $\mu > 1$), but the relative errors start to become very large in the regime of branching caustics (later part of the nonlinear regime). Also, Kasim *et al.* demonstrated the large errors that come with solving the Poisson equation in Kugland *et al.* for a system in the nonlinear regime.

So far, we have seen that existing methods of magnetic field reconstruction require simplifying assumptions in order to get an equation that, when solved, gives the line-integrated transverse magnetic field or magnetic potential, which can be used to derive the magnetic field. This highlights two gaps: (i) In later parts of the nonlinear regime (e.g., branching caustics regime), there is no known reconstruction method despite the fact that nonlinear features do occur in some experimental radiographs [28–30]. In this regime, experimental radiographs are analyzed by comparison to simulated radiographs of a hypothesized magnetic field structure. (ii) In both regimes, there is no reconstruction method that can give the 3D magnetic field. As we will demonstrate in the next few subsections, the proposed artificial neural network method can address both gaps.

B. Artificial neural networks (ANN)

Artificial neural networks are a class of models inspired by biological neural networks, commonly used for tasks that are too complicated for rule-based programming. A neuron, the basic unit of an artificial neural network, takes an input vector \mathbf{x} , and returns a scalar output y given by

$$y = S(\mathbf{w} \cdot \mathbf{x} + b), \tag{5}$$

where \mathbf{w} is the weight vector of the neuron, b is a bias term, and S is an activation (or transfer) function, usually a nonlinear function such as the sigmoid function $s(t) = \frac{1}{1+e^{-t}}$ or the hyperbolic tangent (see Fig. 3).

An artificial neural network consists of many of these neurons joined by the inputs and outputs (i.e., the output of several neurons is fed into the input of another neuron), and parameterized by θ , the vector of weights and biases. There are many variants of artificial neural networks such as recurrent neural networks, cascade-forward networks, and the feedforward neural network. A typical feedforward neural network consists of an input layer with n inputs, hidden layers of arbitrary numbers of neurons and an output layer of m neurons (see Fig. 4). Information only moves forward (hence the name feedforward), and between adjacent layers: The




FIG. 3. A neuron, the basic unit in an artificial neural network. The input vector \mathbf{x} is mapped to a scalar y via a nonlinear function S . The connections represent the inputting of elements of \mathbf{x} into the neuron, and each connection is assigned a weight, which is used in calculating the output.

output from the L th layer strictly goes into the input of the $(L + 1)$ th layer. Essentially, the feedforward neural network is a function that maps a vector in \mathbb{R}^n to a vector in \mathbb{R}^m .

One very powerful feature of feedforward artificial neural networks is that with just a single hidden layer and any appropriate choice of activation function (the sigmoid or hyperbolic tangent suffices), it can approximate any continuous function to any desired degree of accuracy given enough neurons in the hidden layer [31]. As such, feedforward artificial neural networks are universal approximators, so if we encounter errors during application, it must be due to (i) inadequate learning, which can arise from insufficient training data or improper implementation of the training process (more on the training process in the next paragraph), (ii) insufficient neurons in the hidden layer, or (iii) a stochastic relationship between inputs and the desired outputs.

While Ref. [31] established that feedforward artificial neural networks can express any continuous function, we also need to know how to select the appropriate θ so that the artificial neural network is a good approximation to our desired function $\mathbf{g}(\mathbf{x})$. Given data on inputs \mathbf{x} and targets $\mathbf{g}(\mathbf{x})$, we can train the artificial neural network: Upon specifying an objective function (e.g., mean-squared error between the target $\mathbf{g}(\mathbf{x})$ and the output of the artificial neural network $\mathbf{y}(\mathbf{x}, \theta)$), we can use backpropagation [32] (a method of finding the gradient




FIG. 4. Schematic of a feedforward neural network, where the outputs of one layer are fed into the inputs of an adjacent layer. It takes a vector $\mathbf{x} \in \mathbb{R}^n$ and outputs a vector $\mathbf{y}(\mathbf{x}, \theta) \in \mathbb{R}^m$, where θ represents the parameters (weights and biases) of the neural network.

of the objective function with respect to θ) in conjunction with an optimization algorithm (e.g., stochastic gradient descent) to iteratively adjust θ until the objective function is minimized.

Given the merits of artificial neural networks mentioned above, it is no surprise that artificial neural networks are being used in some areas of physics. In astrophysics, artificial neural networks are being used to separate astrophysical signals from the cosmic microwave background [33] and to classify stars and galaxies [34]. In optical tomography, Kamilov *et al.* [25] used artificial neural networks to enhance the beam propagation method [35] and retrieve the index of refraction of a 3D object. The images reconstructed from this technique were found to be of higher quality than images from optical diffraction tomography and Radon tomographic reconstruction. In dynamical electron scattering experiments, Van den Broek *et al.* used artificial neural networks to retrieve 3D potentials [26].

In terms of retrieving information from proton radiographs, artificial neural networks can be used as an input-output map. Given the radiograph pixel values as inputs, the artificial neural network should output useful quantities related to the \mathbf{B} field. Such an artificial neural network can be trained (i.e., the weights are selected) by applying backpropagation and optimization algorithms to many sets of input (radiograph pixels) and target (\mathbf{B} field quantities) examples. After enough training examples, the artificial neural network will become a tool that can estimate \mathbf{B} field quantities given a radiograph. One key strength of the artificial neural network method is that the bulk of the computational cost comes from generating training data and training the artificial neural network, which is a one-off cost.

III. METHODS

A. Reconstruction of an arbitrary \mathbf{B} field

In this section, we outline the steps to using an artificial neural network to reconstruct any arbitrary \mathbf{B} field. First, we expand the magnetic field $\mathbf{B}(\mathbf{r})$ as a linear combination,

$$\begin{aligned} \mathbf{B}(\mathbf{r}) &= \alpha_1 \mathbf{B}_1 \left(\frac{\mathbf{r} - \mathbf{r}_1}{\sigma_1} \right) + \alpha_2 \mathbf{B}_2 \left(\frac{\mathbf{r} - \mathbf{r}_2}{\sigma_2} \right) + \dots \\ &= \sum_{n=1}^N \alpha_n \mathbf{B}_n \left(\frac{\mathbf{r} - \mathbf{r}_n}{\sigma_n} \right), \end{aligned} \quad (6)$$

where N is the number of terms used in the expansion, α_n is a scalar coefficient for the n th term, \mathbf{B}_n is a ‘‘basis’’ magnetic field, \mathbf{r}_n is the position offset of the field, and σ_n is a scaling factor. With enough basis fields, one can describe fields representative of those found in high-energy-density plasmas [18], such as the Weibel instability. While not necessary, these basis fields should be chosen such that most magnetic fields in plasmas can be represented with as few basis fields as possible, so that we require less training data to train the artificial neural network. One possible way to achieve this is to use principal components analysis (PCA) [36] on a large dataset of known \mathbf{B} fields in plasmas. Principal components analysis looks at a large set of multidimensional vectors and first finds the direction of highest variance in the data (the first principal component), and then finds a set of vectors orthogonal to the

first principal component that explains the remainder of the variance. While $\mathbf{B}(\mathbf{r})$ is a vector field, it can be converted into a vector \mathbf{c} for principal components analysis by concatenating the magnetic fields at various different points, e.g., for a grid that runs from 0 to 9 in the x , y , and z directions, we can write

$$\mathbf{c} = \begin{bmatrix} \mathbf{B}(\mathbf{r}_{000}) \\ \mathbf{B}(\mathbf{r}_{001}) \\ \vdots \\ \mathbf{B}(\mathbf{r}_{999}) \end{bmatrix}, \quad (7)$$

where \mathbf{r}_{xyz} is the vector (x, y, z) . If such convenient basis fields cannot be determined, we can use the fact that all magnetic fields can be expanded in a basis written in the form of Eq. (7), where each element of the vector corresponds to a basis field (i.e., \mathbf{B}_1 corresponds to $\mathbf{B}_x(\mathbf{r}_{000})$, \mathbf{B}_2 corresponds to $\mathbf{B}_y(\mathbf{r}_{000})$, \mathbf{B}_3 corresponds to $\mathbf{B}_z(\mathbf{r}_{000})$, \mathbf{B}_4 corresponds to $\mathbf{B}_x(\mathbf{r}_{001})$ and so on).

Next, generate training data by creating variations of the parameters α_n , σ_n , and \mathbf{r}_n in the form

$$\mathbf{g} = \begin{bmatrix} \alpha_1 \\ \mathbf{r}_1 \\ \sigma_1 \\ \vdots \\ \alpha_N \\ \mathbf{r}_N \\ \sigma_N \end{bmatrix}, \quad (8)$$

and then conducting numerical simulations (e.g., using software packages mentioned in Refs. [18,19]) to obtain the radiograph for each variation of the parameters. The radiograph is expressed as a vector where each element represents the intensity of the protons at a specific pixel.

Then, using the radiograph pixel values as inputs and \mathbf{g} as the targets, apply backpropagation and optimization algorithms to train the feedforward neural network [32]. (This procedure selects the weights and biases of the neural network and is done easily in most machine learning frameworks/packages.) After training, the artificial neural network is ready to reconstruct \mathbf{B} fields: Input the radiograph into the artificial neural network to obtain the predicted parameters [in the form of Eq. (8)] and insert these values into Eq. (6). See Fig. 5 for a schematic of the training and reconstruction process.

B. Assumptions, practical considerations, and implementation

In our simulations, we have made some assumptions for simplicity, but these assumptions are not crucial in the success of our approach. We assumed that the probe beam only interacts with the plasma via the \mathbf{B} field (no electric fields or collisions with matter). We also assumed that the proton source is a point source, and the probe beam is a planar sheet (velocities in the z direction, before deflection from the plasma, are uniform). As feedforward artificial neural networks are universal function approximators, in principle the technique outlined in the previous section will still work even if the assumptions are violated (e.g., protons interact with the electric field of the plasma, protons collide with the plasma, proton source is of finite size, probe beam follows a specific angular

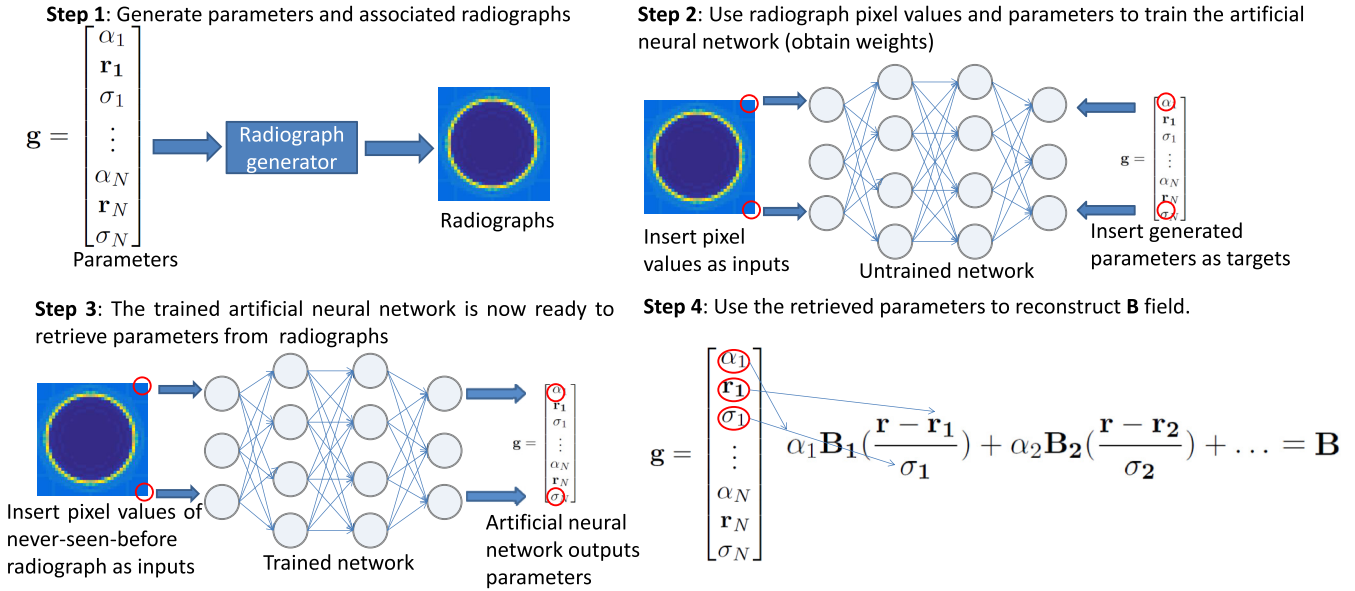


FIG. 5. Schematic of the prediction and reconstruction process.

distribution), as long as we include these effects during the production of training data (radiographs).

To obtain the radiographs we start off with the Lorentz force equation for \mathbf{B} fields only, given by

$$\frac{d\mathbf{v}}{dt} = \frac{e}{m_p} \mathbf{v} \times \mathbf{B}. \quad (9)$$

This equation, along with $\frac{d\mathbf{r}}{dt} = \mathbf{v}$ was numerically integrated given the initial conditions of \mathbf{r} and \mathbf{v} to get the final positions of the protons on the screen. These final coordinates are then binned in order to produce radiographs.

We used a fully connected (dense) feedforward artificial neural network for simplicity, and we will discuss the possibilities of using other types of artificial neural networks in Sec. V. Scaled conjugate gradient was the optimization algorithm of choice during training because: (i) it is not RAM intensive (this is an important factor because in order to get more accurate results, training with more complicated \mathbf{B} fields and higher resolution radiographs are required, resulting in an increase in the number of weights in the artificial neural network. If the optimization algorithm does not scale well, an impractical amount of RAM will be required); (ii) it can take advantage of parallel CPU and GPU computing, allowing it to run effectively on supercomputing clusters.

Before the training process, the entire data set is scaled so that each feature of the input and target (e.g., σ_1 , α_1 , the proton intensity in pixel 1 etc.) falls in the range $[-1, 1]$ to prevent features of small magnitude from converging slowly during optimization [37], and the scaling is undone afterwards. The objective function was chosen to be the mean-squared error (MSE) between the artificial neural network output and the target.

Due to the flexibility of artificial neural networks, overfitting (accidental modeling of noise) is an issue so early stopping and neural network regularization are implemented. In early stopping, training is halted when the errors start increasing on a data set that was not used in the training process [38].

This is done by first splitting the entire simulated data set into training, validation, and testing sets at random in the ratio 70/15/15. The artificial neural network is applied to the training set, and during each iteration the mean-squared error for the validation set is calculated. Initially, after each iteration, the artificial neural network becomes better at modeling the physical phenomenon and the validation mean-squared error will decrease. There will come a point when the artificial neural network starts to model the noise in the training set, and the validation mean-squared error will stop decreasing and eventually start increasing (see Fig. 6 for an illustration). Training is halted after a specified number of iterations fail to decrease the validation mean-squared error. In neural network regularization [39], the objective function is modified by adding a term proportional to the mean-squared weight, and the constant of proportionality is known as the regularization parameter (chosen via cross-validation). This penalizes the




FIG. 6. Typical curves of training and validation errors with respect to training iteration. Beyond a certain point, the artificial neural network starts to model noise, causing the validation error to increase. Training should be halted when this happens.

TABLE I. Parameters for simulations in the following subsections. The proton velocities for subsection F refer to protons in the beam fired in the z direction.

	Value in subsection A–E	Value in subsection F
Ellipsoidal blob parameter a /mm	0.1	0.7
Ellipsoidal blob parameter b /mm		1
Flux rope height/mm	2	0.3
Flux rope parameter a /mm	0.8	0.5
Distance between proton source and object l /mm		7
Distance between object and screen L /mm		93
Number of neurons in input layer	2500	5000
Number of neurons in output layer	2 for subsections A, B, D, F and 1 for subsections C, E	
Velocity of protons in the z direction/ ms^{-1}	10^6	
Velocity of protons in the x, y direction/ ms^{-1}	Ranges from -5×10^4 to 5×10^4	Ranges from -6.9×10^5 to 6.9×10^5

neural network for having large weights or too many neurons, thus encouraging simpler models.

At this point, the artificial neural network is used to predict quantities on the testing set, and the testing errors are indicative of the artificial neural network's overall performance. For each simulated data set, the training process is run with the number of neurons in the hidden layer ranging from 10 to 100 in steps of 10, and the configuration with the lowest value of the objective function (mean-squared error plus regularization term) is initially picked. If the value of the objective function is still decreasing when 100 neurons are used, then the search is extended in steps of 10 till 150 neurons. Once a single-layer configuration is picked, another search is performed with multiple hidden layers (in increments of one layer), up to five hidden layers. Similarly, if the value of the objective function is still decreasing when five hidden layers are used, the search is extended to ten hidden layers. After this search, the configuration with the lowest value of the objective function is picked and reported in the Results section.

C. Retrieval of specific parameters

The idea presented in Sec. III A requires large amounts of data and processing power, and might be more than necessary if the user only intends to retrieve certain parameters of the \mathbf{B} field, such as the peak field strength or the full width half max (FWHM) of a Gaussian magnetic flux rope, instead of reconstructing the entire field. This assumes that the user already knows the remainder of the parameters beforehand. For example, if the user only wants to retrieve the peak \mathbf{B} field strength (proportional to the α_i coefficient), then the model of the \mathbf{B} field, the offset \mathbf{r}_i , and the scaling factor σ_i must be known. In this case, the user can repeat the procedure in Sec. III A, except with the following changes: (i) data is generated by varying only the parameter(s) of interest; (ii) the target vector consists of only the parameter(s) of interest. In fact, this can be applied to parameters other than α_i , \mathbf{r}_i , and σ_i . For example, in an ellipsoidal magnetic blob (which is a spheroid), there are two characteristic lengths, one characterizing the length in the xy plane a and the other characterizing the length along the z -axis b . If the user knows all other parameters and wants to retrieve a and b , then an artificial neural network trained on simulated radiographs which variation is only due to varying values of a and b will do the job.

IV. RESULTS AND DISCUSSION

We have simulated special cases of Eq. (8) as a proof of concept of the idea in Sec. III A. All results shown here come from applying a trained artificial neural network on the testing set (which is not used in training the artificial neural network), and is indicative of the performance when tested on experimental data. Some of the simulation parameters used in the following subsections can be found in Table I.

A. Reconstructing magnetic fields, a proof of concept

Consider the following two fields: (a) a magnetic ellipsoidal blob, representative of fields generated by the Weibel instability [2], that can be written as

$$B_\phi = B_0 \frac{r_0}{a} \exp\left[-\left(\frac{r_0^2}{a^2} + \frac{z_0^2}{b^2}\right)\right], \quad (10)$$

where B_0 is proportional to the peak field strength, r_0 is the distance to the center in the xy plane, z_0 is the distance to the center along the z axis, and a, b are characteristic lengths of the ellipsoid; (b) a magnetic flux rope of Gaussian cross section, representative of fields due to laser generated plasma flows [40], can be written as

$$B_y = B_0 \exp\left(-\frac{x_0^2 + z_0^2}{a^2}\right), \quad (11)$$

where B_0 is the peak field strength, x_0 and z_0 are the distances to the center along the x and z axes, respectively, and a is a characteristic length of the Gaussian.

In terms of Eq. (8), we assign α_1 to B_0 of the magnetic ellipsoidal blob and α_2 to B_0 of the magnetic flux rope. α_1 was varied from 5 to 6 T (defocusing) in steps of 0.01 T while α_2 was varied from 2.01 to 3 T in steps of 0.03 T, and all other parameters were kept constant. Radiographs of 50 by 50 pixels were generated for each configuration. As mentioned earlier, 70% of these radiographs were randomly chosen to train the artificial neural network, 15% were randomly assigned to the validation set to prevent overfitting, and the trained artificial neural network was used to predict the α_1 and α_2 values on the remainder 15% of the radiographs. The errors, defined as $|\frac{\text{predicted value} - \text{actual value}}{\text{actual value}}|$, are plotted in Fig. 7. We see that nearly all the errors are less than 5%, suggesting that the full scale implementation outlined in Sec. III A will work given enough basis fields. Though there are some undesirable outliers in α_2 ,




FIG. 7. Error histograms for the α coefficients of two basis fields, using an artificial neural network with 1 hidden layer consisting of 10 neurons. The mean errors are 0.34% and 2.74%, while the median errors are 0.20% and 1.29% for α_1 (ellipsoidal blob, left) and α_2 (flux rope, right), respectively. More parameters can be found in Table I.




FIG. 9. Error histograms for α (left) and σ (right) for a magnetic ellipsoidal blob, using an artificial neural network with 1 hidden layer consisting of 50 neurons. The mean errors are 0.26% and 0.05% while the median errors are 0.20% and 0.04% for α and σ , respectively. More parameters can be found in Table I.

it is likely to be a result of inadequate data rather than a flaw in the artificial neural network method. This will be discussed in Sec. IV E.

B. Obtaining B field parameters from a magnetic ellipsoidal blob

In this subsection, we demonstrate that (i) the artificial neural network method works in the nonlinear regime, and (ii) the parameter retrieval concept in Sec. III C can be done. Here, we retrieve the field strength coefficient α and the scaling factor σ .

Radiographs for a magnetic ellipsoidal blob were generated with α [representing B_0 in Eq. (10)] ranging from 0.1 to 0.3 T (defocusing) in steps of 2×10^{-4} T and σ ranging from 0.9 to 1 in steps of 0.02. This spectrum of α spans both the caustic and noncaustic regime, as can be seen by the radiographs plotted in Fig. 8. The histogram of errors are plotted in Fig. 9. We can see that the average errors are well below 1%, suggesting that artificial neural networks can be used for parameter retrieval,




FIG. 8. Left: Radiograph for a magnetic ellipsoidal blob at $B = 0.1$ T, $\sigma = 1$. This is in the noncaustic regime, where the ring around the center is smeared out. Right: Radiograph for a magnetic ellipsoidal blob at $B = 0.3$ T, $\sigma = 1$. This is in the caustic regime, where most of the protons fall into a very thin ring. The scales are in arbitrary units.

an alternative to reconstructing entire magnetic fields. We also see that this method works in the nonlinear regime, where $\mu \approx 2$.

C. Branching caustics

The power diagram method [22] gives relative errors of more than 10% in the branching caustics regime. Here we show that the artificial neural network method is flexible enough to accommodate this scenario. We extend the range of field strengths in Sec. IV B to range from 0.1 to 1.5 T in steps of 2×10^{-4} T, spanning the linear, caustic, and branching caustic regime. As an illustration, the horizontal profile of the radiograph at 1.5 T (branching caustics regime) is plotted in Fig. 10. The error histogram is plotted in Fig. 11 and we can see that all the errors are well below 1%.

D. Effect of noise on accuracy

So far we have shown that artificial neural networks trained on noise-free radiographs can retrieve quantities from noise-free radiographs with a high accuracy. We proceed to explore the changes in accuracy when noise is introduced into all the radiographs (training, validation, and testing sets). Suppose a pixel in the radiograph has a value of χ and we want to introduce random noise of $x\%$. Then each pixel is replaced by a random value from a Gaussian distribution with a mean of χ and a standard deviation of $\chi \times x\%$. This was done for $x = 5, 10, 20,$ and 30 percent on the radiographs in Sec. IV B, and the entire process of training and prediction was repeated. The error histograms for α and σ are plotted in Fig. 12. We notice that for both quantities, it takes an increase from 5% to 30% noise in order to roughly double the mean and median errors. This demonstrates the robustness of artificial neural networks to noise, although noise does occasionally cause very high errors. It is worth noting that the right model to




FIG. 10. Horizontal profile of a radiograph for a magnetic ellipsoidal blob at 1.5 T, the branching caustics regime. Notice that there are two maxima in the intensity profile, instead of one in the case of the caustic regime.

use is Poisson noise, but that model approximates Gaussian noise for a large number of particles per pixel, which is true in our case. The relationship between errors and input noise for various configurations of artificial neural networks is further explored in Refs. [41–44].

E. Effect of the amount of training data on accuracy

While the artificial neural network method seems promising so far, it is reliant on the large amounts of training data (specifically, the amount of information in the data, or information entropy) for its accuracy. To elucidate this fact, we generated




FIG. 11. Error histogram of α for a magnetic ellipsoidal blob spanning the linear, caustic, and branching caustic regime, using an artificial neural network with 1 hidden layer consisting of 10 neurons. The mean error is 0.06% while the median error is 0.04%. More parameters can be found in Table I.

data for a magnetic ellipsoidal blob and varied only α between the values 0.1 to 1.5 T, similar to the scenario in Sec. IV C, except with a larger step size of 10^{-3} T. The error histogram is plotted in Fig. 13. In comparison to Fig. 11, we see that having a larger step size and thus having less information leads to an increase in errors. This, combined with the universality of the artificial neural network proved in Ref. [31], suggests that extreme outliers in errors can be overcome by generating more data that increases the information entropy of the data set and retraining the neural network. The relationship between errors and size of data set for various configurations of artificial neural networks is further explored in Refs. [44–47].

F. Limitations of proton radiography and the need for proton tomography

Proton radiographs do not necessarily form one-to-one relationships with field structures: Suppose that at the edge of a plasma that is facing the proton beam, there is a very strong \mathbf{B} field that deflects the incoming protons before these protons could penetrate further. Then the radiograph formed is independent of the \mathbf{B} fields in the remainder of the plasma, because no protons interact with it. Due to the lack of information in the radiographs, no method can fully reconstruct the \mathbf{B} fields. As such, there is a need to modify the experimental setup to capture more information from the \mathbf{B} field.

One possible way to capture more information is to include more probe beams in different directions (tomography). As an example, consider two adjacent field structures, the ellipsoidal blob (field strength parameter assigned to α_1 , ranging from 9 to 9.25 T in steps of 0.005 T) and the flux rope (field strength parameter assigned to α_2 , ranging from 0.3 to 0.4 T in steps of 0.002 T), with the former obscuring the latter in the z direction by 0.5 mm. When the artificial neural network method was used on radiographs due to a beam fired in the z direction, the errors in the retrieved field strength of the flux rope are very high (top panel, Fig. 14) due to the lack of protons probing the field structure. When another probe beam was fired in the x direction (with x velocity of 10^6 ms^{-1} and y, z velocities ranging from $-2 \times 10^5 \text{ ms}^{-1}$ to $2 \times 10^5 \text{ ms}^{-1}$) and the radiographs were used in addition to the ones from the probe beam in the z direction, errors for both field strengths decrease by at least an order of magnitude (bottom panel, Fig. 14). This demonstrates two facts: (i) the artificial neural network method (and any other method) cannot fully reconstruct magnetic fields if the radiographs carry insufficient information; (ii) including more information decreases errors, even if the field structure is not obscured, as can be seen by the reduction in errors for α_1 in Fig. 14. We hope this will inspire future work on theoretical error bounds in artificial neural networks given the lack of information in the data set.

G. Comparison with the existing radiograph inversion techniques

The artificial neural network method addresses the two gaps in existing reconstruction techniques, by being able to work in the nonlinear regime (such as the branching caustic regime) and to produce 3D reconstruction of the magnetic field. While existing inversion techniques rely on the paraxial limit for

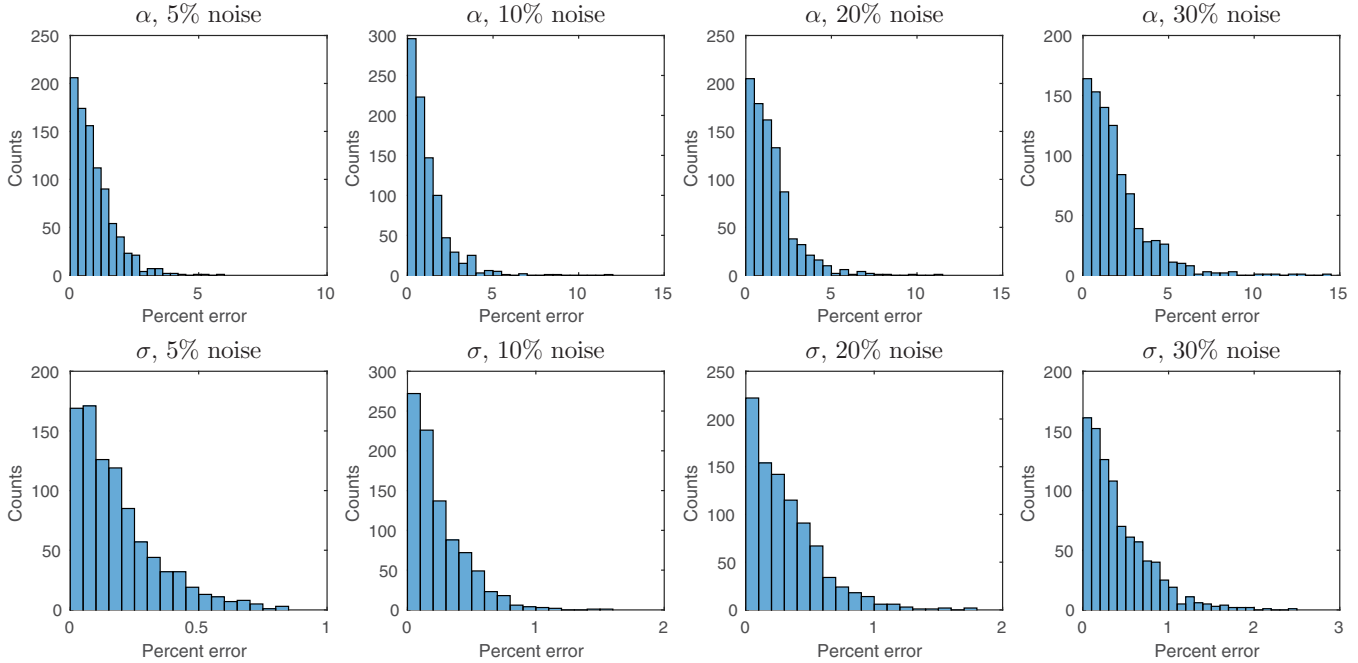


FIG. 12. Error histograms of α and σ when 5, 10, 20, and 30 percent noise is introduced into the radiographs for a magnetic ellipsoidal blob. The artificial neural network configurations are: 8 hidden layers with 80 neurons per layer, 5 hidden layers with 50 neurons per layer, 6 hidden layers with 70 neurons per layer, and 7 hidden layers with 100 neurons per layer for 5, 10, 20, and 30 percent noise, respectively. For α , the mean errors are 0.92%, 1.14%, 1.49%, and 1.91%, while the median errors are 0.73%, 0.82%, 1.19%, and 1.46% for 5, 10, 20, and 30 percent noise, respectively. For σ , the mean errors are 0.18%, 0.24%, 0.31%, and 0.41%, while the median errors are 0.14%, 0.18%, 0.25%, and 0.31% for 5, 10, 20, and 30 percent noise, respectively. Notice that it takes an increase from 5% noise to 30% noise in order to roughly double the mean and median errors, indicating that the artificial neural network method is robust to noise. More parameters can be found in Table I.

simplicity, the artificial neural network technique does not rely on such a limit, and in fact would benefit more if the paraxial limit was not used—the protons should ideally have nonzero velocities in the x and y directions so that it will




FIG. 13. Error histogram in α for a magnetic ellipsoidal blob (more parameters in Table I) when the step size is increased by a factor of 5, leading to less data. The mean error is 0.20% and the median error is 0.16%, using an artificial neural network with 7 hidden layers with 40 neurons per layer.

be deflected by \mathbf{B}_z , allowing the artificial neural network to capture more information and thus reconstruct the magnetic fields more accurately.

Also, existing techniques assume that the protons move in a straight line within the plasma, but this assumption does not hold when the \mathbf{B} field is so strong that deflection occurs within the plasma. As a result, the existing techniques will inevitably fail in the limit of extremely large \mathbf{B} fields. In comparison, the neural network method will work because it does not require this assumption.

One major benefit of using artificial neural networks is the long-run computational cost savings. Generating each set of 50 by 50 pixel radiographs (one radiograph for each variation of parameters) takes on the order of hours and days using 16 cores on one node of the Arcus Phase B supercomputer [48]. Training the artificial neural network takes on the order of minutes and hours when using a single GPU on the Arcus Phase B, for neural networks with up to 10 hidden layers, with each layer consisting of up to 150 neurons. Reconstruction takes on the order of seconds without using any parallel processing or GPU. If this project were to go full-scale, we can see that most of the computational cost is in the generation of training data and the training of the artificial neural network, which is a one-off cost. In comparison, existing methods of reconstruction have a recurring cost. As such, over the long run, if the artificial neural network is used to invert sufficiently many radiographs, the artificial neural network method is computationally more efficient.

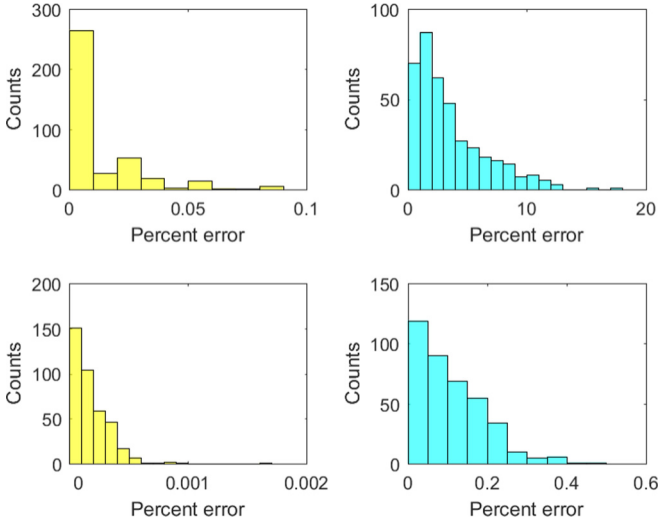


FIG. 14. Error histograms of α_1 (left) and α_2 (right) for two scenarios. Top panels: Using radiographs generated by proton beams in the z direction as training data, the mean errors are $1.20 \times 10^{-2}\%$ and 3.48% while the median errors are $3.59 \times 10^{-3}\%$ and 2.50% for α_1 and α_2 , respectively, using an artificial neural network with 1 hidden layer consisting of 30 neurons (more parameters in Table I). The errors for α_2 (field strength of the flux rope) are high because the field structure associated with α_1 (ellipsoidal blob) is deflecting many protons away from the flux rope, causing a lack of information in the resulting radiographs. Bottom panels: Using radiographs generated by proton beams in the z and x directions as training data, the mean errors are $1.81 \times 10^{-4}\%$ and 0.11% while the median errors are $1.41 \times 10^{-4}\%$ and $8.68 \times 10^{-2}\%$ for α_1 and α_2 , respectively, using an artificial neural network with 1 hidden layer consisting of 110 neurons. We see that upon including data from the proton beam in the x direction, more information for both field structures is added to the data set and errors reduce by at least an order of magnitude.

However, the artificial neural network method has some drawbacks. For example, the overall accuracy of the artificial neural network can only be determined empirically, whereas error propagation can be performed for existing techniques. While the artificial neural network method will allow for computational cost savings over the long run, the minimal startup computational cost to get it working for a nontrivial field structure is quite high, because the artificial neural network must be trained with many basis fields before it can be used. This is in contrast to existing techniques, where any field, as long as the assumptions are met, can be imaged with the computational cost of solving a differential equation.

V. CONCLUSIONS AND FUTURE WORK

In conclusion, we have reviewed existing techniques on analyzing \mathbf{B} fields from proton radiography and the basics of artificial neural networks. Using the fact that artificial neural networks are highly flexible function approximators, we proposed for the first time the idea of using artificial neural networks to reconstruct arbitrary \mathbf{B} fields and retrieve important field parameters.

Via simulations, we showed that an artificial neural network can reconstruct \mathbf{B} fields that can be expressed as linear

combinations of two fields, and retrieve useful quantities of \mathbf{B} fields such as characteristic lengths. We also explored the effects of noise and size of data set on the accuracy of the artificial neural network, and found that artificial neural networks are robust to noise. Artificial neural networks can accommodate a wide variety of scenarios and assumptions which existing techniques cannot, such as the branching caustics part of the nonlinear regime. We also highlighted the need for proton tomography as certain field structures cannot be reconstructed fully due to the lack of information from a single radiograph.

As the usage of artificial neural networks in diagnosing \mathbf{B} fields in high-energy-density plasmas is new, there are many avenues where this work can be developed further. For example, we can explore the effects of a mismatch in noise between the training set and testing set, because the noise in experimental data is not known beforehand. There are at least three ways to improve the accuracy of the artificial neural network: (i) experiment with other types of artificial neural network architecture. For example, convolutional neural networks are a type of feedforward artificial neural network where the connections between neurons are inspired by the animal visual cortex [49] and, as such, perform very well in image recognition. Since radiograph inversion involves image recognition, convolutional neural networks could offer better performance than the fully connected (dense) feedforward neural network used in this paper. Recurrent neural networks are a class of artificial neural networks where the neuron connections form directed cycles, and such architecture has advantages in analyzing time series data. We could use recurrent neural networks on a time series of proton radiographs to shed light on the dynamics of the \mathbf{B} field and hence the plasma. (ii) Include energy-resolved radiographs. In our simulations, we only looked at the spatial distribution of the protons, so including extra information on the proton energies could improve accuracy. (iii) Study the effects of the number of pixels on accuracy. It is interesting to note that promising results were obtained despite the low resolution of the radiographs (50 by 50 pixels). Understanding the effects of discretization noise could help us determine the quality of radiographs to be generated in order to train an artificial neural network to a specific accuracy.

Furthermore, the artificial neural network approach can be extended to similar systems, such as diagnosing electric fields in plasmas or characterizing micromagnetic patterns in magnetic media via electron scattering [50]. Finally, a full-scale implementation of an artificial neural network that can reconstruct any \mathbf{B} field is a possibility we can look forward to.

ACKNOWLEDGMENTS

The authors acknowledge the support from the plasma physics HEC Consortium EPSRC Grant No. EP/L000237/1, as well as the Hartree Centre, Daresbury Laboratory, Central Laser Facility, and the Scientific Computing Department at the Rutherford Appleton Laboratory. The authors acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work [48]. N.C. acknowledges financial support from the Singapore government. M.F.K. gratefully thanks the Indonesian Endowment

Fund for Education for its support. M.C.L. thanks the Royal Society Newton International Fellowship for support. The

authors acknowledge support from OxCHEDS and P.A.N. for his William Penney Fellowship with AWE plc.

-
- [1] S. Eliezer, *Plasma Phys. Control. Fusion* **45**, 181 (2003).
- [2] C. Huntington, F. Fiuza, J. Ross, A. Zylstra, R. Drake, D. Froula, G. Gregori, N. Kugland, C. Kuranz, M. Levy *et al.*, *Nature Phys.* **11**, 173 (2015).
- [3] M. L. Bernet, F. Miniati, S. J. Lilly, P. P. Kronberg, and M. Dessauges-Zavadsky, *Nature* **454**, 302 (2008).
- [4] O. V. Gotchev, P. Y. Chang, J. P. Knauer, D. D. Meyerhofer, O. Polomarov, J. Frenje, C. K. Li, M. J.-E. Manuel, R. D. Petrasso, J. R. Rygg, F. H. Séguin, and R. Betti, *Phys. Rev. Lett.* **103**, 215004 (2009).
- [5] M. Basko, A. Kemp, and J. Meyer-ter-Vehn, *Nuclear Fusion* **40**, 59 (2000).
- [6] P. M. Nilson, L. Willingale, M. C. Kaluza, C. Kamperidis, S. Minardi, M. S. Wei, P. Fernandes, M. Notley, S. Bandyopadhyay, M. Sherlock, R. J. Kingham, M. Tatarakis, Z. Najmudin, W. Rozmus, R. G. Evans, M. G. Haines, A. E. Dangor, and K. Krushelnick, *Phys. Rev. Lett.* **97**, 255001 (2006).
- [7] M. Roth, *J. Instrument.* **6**, R09001 (2011).
- [8] National Research Council, Division on Engineering and Physical Sciences, Board on Physics and Astronomy, Committee on High Energy Density Plasma Physics, and Plasma Science Committee, *Frontiers in High Energy Density Physics: The X-Games of Contemporary Science* (National Academies Press, 2003).
- [9] A. Ravasio, L. Romagnani, S. Le Pape, A. Benuzzi-Mounaix, C. Cecchetti, D. Batani, T. Boehly, M. Borghesi, R. Dezulian, L. Gremillet, E. Henry, D. Hicks, B. Loupias, A. MacKinnon, N. Ozaki, H. S. Park, P. Patel, A. Schiavi, T. Vinci, R. Clarke, M. Notley, S. Bandyopadhyay, and M. Koenig, *Phys. Rev. E* **82**, 016407 (2010).
- [10] F. Nürnberg, M. Schollmeier, E. Brambrink, A. Blažević, D. Carroll, K. Flippo, D. Gautier, M. Geissel, K. Harres, B. Hegelich *et al.*, *Rev. Sci. Instrum.* **80**, 033301 (2009).
- [11] S. Kar, M. Borghesi, P. Audebert, A. Benuzzi-Mounaix, T. Boehly, D. Hicks, M. Koenig, K. Lancaster, S. Lepape, A. Mackinnon, P. Norreys, P. Patel, and L. Romagnani, *High Energy Density Phys.* **4**, 26 (2008).
- [12] C. K. Li, F. H. Séguin, J. A. Frenje, J. R. Rygg, R. D. Petrasso, R. P. J. Town, P. A. Amendt, S. P. Hatchett, O. L. Landen, A. J. Mackinnon, P. K. Patel, V. A. Smalyuk, T. C. Sangster, and J. P. Knauer, *Phys. Rev. Lett.* **97**, 135003 (2006).
- [13] L. Willingale, A. G. R. Thomas, P. M. Nilson, M. C. Kaluza, S. Bandyopadhyay, A. E. Dangor, R. G. Evans, P. Fernandes, M. G. Haines, C. Kamperidis, R. J. Kingham, S. Minardi, M. Notley, C. P. Ridgers, W. Rozmus, M. Sherlock, M. Tatarakis, M. S. Wei, Z. Najmudin, and K. Krushelnick, *Phys. Rev. Lett.* **105**, 095001 (2010).
- [14] C. K. Li, F. H. Sguin, J. A. Frenje, M. Rosenberg, A. B. Zylstra, R. D. Petrasso, P. A. Amendt, J. A. Koch, O. L. Landen, H. S. Park, H. F. Robey, R. P. J. Town, A. Casner, F. Philippe, R. Betti, J. P. Knauer, D. D. Meyerhofer, C. A. Back, J. D. Kilkenny, and A. Nikroo, *Plasma Phys. Control. Fusion* **52**, 124027 (2010).
- [15] C. K. Li, F. H. Séguin, J. A. Frenje, M. J. Rosenberg, H. G. Rinderknecht, A. B. Zylstra, R. D. Petrasso, P. A. Amendt, O. L. Landen, A. J. Mackinnon, R. P. J. Town, S. C. Wilks, R. Betti, D. D. Meyerhofer, J. M. Soures, J. Hund, J. D. Kilkenny, and A. Nikroo, *Phys. Rev. Lett.* **108**, 025001 (2012).
- [16] M. Borghesi, G. Sarri, C. Cecchetti, I. Kourakis, D. Hoarty, R. Stevenson, S. James, C. Brown, P. Hobbs, J. Lockyear, J. Morton, O. Willi, R. Jung, and M. Dieckmann, *Laser Particle Beams* **28**, 277 (2010).
- [17] L. Willingale, P. M. Nilson, A. G. R. Thomas, J. Cobble, R. S. Craxton, A. Maksimchuk, P. A. Norreys, T. C. Sangster, R. H. H. Scott, C. Stoeckl, C. Zulick, and K. Krushelnick, *Phys. Rev. Lett.* **106**, 105002 (2011).
- [18] M. C. Levy, D. D. Ryutov, S. C. Wilks, J. S. Ross, C. M. Huntington, F. Fiuza, D. A. Martinez, N. L. Kugland, M. G. Baring, and H.-S. Park, *Rev. Sci. Instrum.* **86**, 033302 (2015).
- [19] S. Kar, M. Borghesi, C. A. Cecchetti, L. Romagnani, F. Ceccherini, T. V. Liseykina, A. Macchi, R. Jung, J. Osterholz, O. Willi, L. A. Gizzi, A. Schiavi, M. Galimberti, and R. Heathcote, *New J. Phys.* **9**, 402 (2007).
- [20] N. L. Kugland, D. D. Ryutov, C. Plechaty, J. S. Ross, and H.-S. Park, *Rev. Sci. Instrum.* **83**, 101301 (2012).
- [21] C. Graziani, P. Tzeferacos, D. Q. Lamb, and C. Li, [arXiv:1603.08617](https://arxiv.org/abs/1603.08617) (2016).
- [22] M. F. Kasim, L. Ceurvorst, N. Ratan, J. Sadler, N. Chen, A. Savert, R. Trines, R. Bingham, P. N. Burrows, M. C. Kaluza *et al.*, *Phys. Rev. E* **95**, 023306 (2017).
- [23] Y. LeCun, Y. Bengio, and G. Hinton, *Nature* **521**, 436 (2015).
- [24] T. Ciodaro, D. Deva, J. M. de Seixas, and D. Damazio, *J. Phys.: Conf. Ser.* **368**, 012030 (2012).
- [25] U. S. Kamilov, I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, and D. Psaltis, *Optica* **2**, 517 (2015).
- [26] W. Van den Broek and C. T. Koch, *Phys. Rev. Lett.* **109**, 245502 (2012).
- [27] T. E. Cowan, J. Fuchs, H. Ruhl, A. Kemp, P. Audebert, M. Roth, R. Stephens, I. Barton, A. Blazevic, E. Brambrink, J. Cobble, J. Fernández, J.-C. Gauthier, M. Geissel, M. Hegelich, J. Kaae, S. Karsch, G. P. Le Sage, S. Letzring, M. Manclossi, S. Meyeroneinc, A. Newkirk, H. Pépin, and N. Renard-LeGalloudec, *Phys. Rev. Lett.* **92**, 204801 (2004).
- [28] S. N. Chen, E. d'Humières, E. Lefebvre, L. Romagnani, T. Toncian, P. Antici, P. Audebert, E. Brambrink, C. A. Cecchetti, T. Kudyakov, A. Pipahl, Y. Sentoku, M. Borghesi, O. Willi, and J. Fuchs, *Phys. Rev. Lett.* **108**, 055001 (2012).
- [29] N. Kugland, D. Ryutov, P. Chang, R. Drake, G. Fiksel, D. Froula, S. Glenzer, G. Gregori, M. Grosskopf, M. Koenig *et al.*, *Nature Phys.* **8**, 809 (2012).
- [30] H.-S. Park, C. M. Huntington, F. Fiuza, R. P. Drake, D. H. Froula, G. Gregori, M. Koenig, N. L. Kugland, C. C. Kuranz, D. Q. Lamb, M. C. Levy, C. K. Li, J. Meinecke, T. Morita, R. D. Petrasso, B. B. Pollock, B. A. Remington, H. G. Rinderknecht, M. Rosenberg, J. S. Ross, D. D. Ryutov, Y. Sakawa, A. Spitkovsky, H. Takabe, D. P. Turnbull, P. Tzeferacos, S. V. Weber, and A. B. Zylstra, *Phys. Plasmas* **22**, 056311 (2015).

- [31] K. Hornik, M. Stinchcombe, and H. White, *Neural Networks* **2**, 359 (1989).
- [32] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Nature* **323**, 533 (1986).
- [33] C. Baccigalupi, L. Bedini, C. Burigana, G. De Zotti, A. Farusi, D. Maino, M. Maris, F. Perrotta, E. Salerno, L. Toffolatti *et al.*, *Monthly Notices Roy. Astron. Soc.* **318**, 769 (2000).
- [34] S. Andreon, G. Gargiulo, G. Longo, R. Tagliaferri, and N. Capuano, *Monthly Notices Roy. Astron. Soc.* **319**, 700 (2000).
- [35] J. V. Roey, J. van der Donk, and P. E. Lagasse, *J. Opt. Soc. Am.* **71**, 803 (1981).
- [36] K. Pearson, *Philos. Mag. Series 6* **2**, 559 (1901).
- [37] Y. LeCun, L. Bottou, G. B. Orr, and K. R. Müller, Efficient backprop, in *Neural Networks: Tricks of the Trade*, edited by G. B. Orr and K.-R. Müller (Springer, Berlin/Heidelberg, 1998), pp. 9–50.
- [38] R. Caruana, S. Lawrence, and C. L. Giles, *Advances in Neural Information Processing Systems 13: Proceedings of the 2000 Conference* (Cambridge, Boston, 2001), Vol. 13, p. 402.
- [39] A. Krogh, and J. A. Hertz, *Advances in Neural Information Processing Systems* (Morgan Kaufmann, San Francisco, CA, 1992), pp. 950–957.
- [40] S. K. P. Tripathi and W. Gekelman, *Phys. Rev. Lett.* **105**, 075005 (2010).
- [41] J. L. Bernier, A. F. Díaz, F. Fernández, A. Cañas, J. González, P. Martin-Smith, and J. Ortega, *Neural Process. Lett.* **18**, 35 (2003).
- [42] N. W. Townsend and L. Tarassenko, *IEEE Trans. Neural Networks* **10**, 217 (1999).
- [43] J. L. Bernier, J. Ortega, E. Ros, I. Rojas, and A. Prieto, *Neural Comput.* **12**, 2941 (2000).
- [44] K. Hagiwara, T. Hayasaka, N. Toda, S. Usui, and K. Kuno, *Neural Networks* **14**, 1419 (2001).
- [45] P. Niyogi and F. Girosi, *Neural Comput.* **8**, 819 (1996).
- [46] P. Niyogi and F. Girosi, *Adv. Comput. Math.* **10**, 51 (1999).
- [47] A. R. Barron, *Mach. Learn.* **14**, 115 (1994).
- [48] A. Richards, University of Oxford Advanced Research Computing (2015), <http://dx.doi.org/10.5281/zenodo.22558>.
- [49] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, *Proc. IEEE* **86**, 2278 (1998).
- [50] M. Mankos, M. R. Scheinfein, and J. M. Cowley, *IEEE Trans. Magnet.* **32**, 4150 (1996).