

Machine learning assisted analysis of visible spectroscopy in pulsed-power-driven plasmas

Rishabh Datta, Faez Ahmed, and Jack D. Hare

Abstract—We use machine learning models to predict ion density and electron temperature from visible emission spectra, in a high energy density pulsed-power-driven aluminum plasma, generated by an exploding wire array. Radiation transport simulations, which use spectral emissivity and opacity values generated using the collisional-radiative code PrismSPECT, are used to determine the spectral intensity generated by the plasma along the spectrometer’s line of sight. The spectra exhibit Al-II and Al-III lines, whose line ratios and line widths vary with the density and temperature of the plasma. These calculations provide a 2500-size synthetic dataset of 400-dimensional intensity spectra, which is used to train and compare the performance of multiple machine learning models on a 3-variable regression task. The AutoGluon model performs best, with an R2-score of roughly 98% for density and temperature predictions. Simpler models (random forest, k-nearest neighbor, and deep neural network) also exhibit high R2-scores ($> 90\%$) for density and temperature predictions. These results demonstrate the potential of machine learning in providing rapid or real-time analysis of emission spectroscopy data in pulsed-power-driven plasmas.

I. INTRODUCTION

Spectroscopy is a powerful technique for inferring plasma parameters from emitted electromagnetic radiation. For instance, line widths and line ratios can be used to determine electron density and temperature [1]–[3], velocity can be determined from the Doppler shift of spectral lines [1], [4], and magnetic field strength can be inferred from the Zeeman splitting of line radiation [5], [6]. The wide applicability of spectroscopy makes it an attractive tool for implementation in a variety of laboratory plasmas [1], [4], [7]–[9].

In emission spectroscopy, a typical intensity spectrum can contain several peaks (called emission lines) overlaid on a continuum [1]. The lines correspond to bound-bound electron transitions in the ions of the plasma, while the continuum emission results from free-free (Bremsstrahlung emission) and free-bound electron transitions (recombination radiation) [1], [4]. Line radiation generated by the plasma arises due to either collisional or radiative processes [1]. Collisional processes, such as electron impact excitation/de-excitation and three-body recombination, change the energy levels of bound electrons via collisions with other electrons [1], [4]. Similarly, radiative processes, such as photoexcitation/de-excitation, induce energy transitions due to the interaction of bound electrons with photons [1], [4]. Collisional-radiative models balance the rates of excitation (and ionization) against that of de-excitation (and deionization), to determine the spectral emissivity and opacity of radiation emitted from the plasma [2].

The typical approach to determining the ion density n_i from emission spectra is to identify lines dominated by Stark (collisional) broadening, and then to compare the line widths with tabulated data [1], [9], [10], or with the predictions of CR codes, such as PrismSPECT [11], [12]. Similarly, for the characterization of electron temperature T_e , we typically compare the intensity ratios of two or more lines (typically, inter-stage lines for which density changes have a small effect) with the predictions of CR models [1], [3], [10].

When the plasma is not optically thin, radiation transport, which describes how the energy distribution of radiation changes as it propagates through an absorbing, emitting, and/or scattering medium, must be adequately modeled for accurate interpretation of emission spectra. The optical thickness of a material to radiation of frequency ω is characterized by $\tau \equiv \int \alpha(\omega, s) ds$, which is the line-integral of the spectral opacity $\alpha(\omega)$ along the path s [4], [13]. When $\tau \ll 1$, the plasma is optically thin, and the output spectrum is simply the line-integrated emissivity $\epsilon(\omega)$ of the plasma along the path s . Similarly, for $\tau \gg 1$, the plasma is optically thick, and the output spectrum (for a plasma in local thermodynamic equilibrium) is Planckian [4], [13]. In plasmas that are not optically thin, the radiation spectrum recorded by the spectrometer is significantly altered by radiation transport. Furthermore, if the plasma exhibits spatial inhomogeneity along the line-of-sight (LOS), the resulting spectrum may be dominated by strongly-emitting or absorbing regions.

In high-energy-density pulsed-power-driven systems, the condition for optical thinness may not be satisfied [14], [15]. In this paper, our focus is to diagnose pulsed-power-driven plasmas for laboratory astrophysics applications, which typically exhibit ion densities and electron temperatures between $n_i \approx 1 \times 10^{17} - 1 \times 10^{19} \text{ cm}^{-3}$ and $T_e \approx 1 - 50 \text{ eV}$ respectively [14], [16]–[18]. Pulsed-power devices generate plasma flows by driving large currents (1-30 MA) through thin $\sim 10 - 100 \mu\text{m}$ diameter wires. These plasmas are not optically thin to visible radiation, therefore radiation transport modeling becomes important for spectral analysis.

A key drawback of the aforementioned emission spectroscopy analysis approach is that it requires significant CR and radiation transport modeling for the analysis of a given spectrum. The use of machine learning (ML) models can reduce the computational time required for spectral analysis, especially for large batches of spectral data, and provide rapid real-time results during experimentation. Spectroscopy has previously been combined with supervised ML techniques, primarily in lower-density plasmas. Visible emission spectroscopy combined with regression methods and neural networks has been used to predict density and temperature in

R. Datta and J. D. Hare are with the Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, USA

F. Ahmed is with the Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, USA

low-temperature low-density ($n_e \sim 1 \times 10^{10} \text{ cm}^{-3}$, $T_e \sim 1 \text{ eV}$) laboratory plasmas [19]. Similarly, neural networks have been used to predict electron energy distribution in low-temperature non-thermal plasmas [20], and classifiers have been used for trace element and impurity detection in RF-generated plasmas [21]. Neural network regressors and classifiers have also been shown to accurately predict electron temperature and divertor detachment from UV/XUV spectroscopy measurements in magnetic confinement fusion devices [22]–[24]. In the examples above, the labeled dataset for training is generated by simultaneous spectroscopy and independent density and temperature measurements with other diagnostics, like Thompson scattering or Langmuir probes [22]–[24]. This eliminates the need for complicated theoretical or computational modeling. However, simultaneous independent measurements using secondary diagnostics are not always possible; for example, in experiments with poor diagnostic access or diagnostic unavailability. Moreover, alternative diagnostics such as Langmuir probes also perturb the plasma. Lastly, in experiments with low repetition rates, such as in pulsed-power plasmas, it is challenging to generate a purely experimental dataset for the data-hungry supervised ML task. In such situations, we must therefore rely on synthetic data for training.

In this paper, we use a machine learning approach to predict electron density and temperature profiles along the spectrometer line of sight (LOS) from visible emission spectroscopy data in a pulsed-power-driven aluminum plasma. In contrast to previous work [19], [20], [23], [24], which focuses on lower-density and/or optically thin plasmas, here we aim to characterize high-density non-optically thin spatially-inhomogeneous plasmas, characteristic of pulsed-power-driven plasmas generated using exploding wire arrays [16], [17]. We approach the problem in two parallel ways. In the first approach, we frame the prediction problem as a single-objective optimization problem, where we minimize the deviation between a simulated spectrum and a target spectrum. We generate the simulated spectra using CR calculations performed using PrismSPECT, which are then used to solve radiation transport in a spatially-inhomogeneous plasma. In the second approach, we solve a multi-variable regression problem, where we predict ion density $n_i(s)$ and electron temperature $T_e(s)$ as a function of position s along the spectrometer LOS from a given spectrum. We train multiple supervised ML models — linear regressor, k-nearest neighbor, decision trees, random forest, deep and convolution neural networks, and AutoGluon — using synthetic data, generated from radiation transport simulations. The AutoGluon model performs best, with an R2-score of roughly 98% for density and temperature predictions, and significantly reduces the computation time over optimization-based curve fitting methods. Our results demonstrate the potential of machine learning methods in providing rapid or real-time analysis of emission spectroscopy data in pulsed-power-driven plasmas.

II. DATASET GENERATION

A. Radiation Transport Modeling

We use PrismSPECT [12] to compute emissivity $\epsilon(\omega)$ and opacity $\alpha(\omega)$ values for an aluminum plasma, in the visible




Fig. 1: Comparison of the output intensity calculated by the radiation transport solver with that from a planar plasma simulation in PrismSPECT ($n_i = 5 \times 10^{17} \text{ cm}^{-3}$, $T_e = 2.5 \text{ eV}$, and $L = 10 \text{ mm}$), and with a zero opacity $\tau = 0$ case.

range of the electromagnetic spectrum ($400 \text{ nm} < \lambda < 700 \text{ nm}$). We use a steady-state non local thermodynamic equilibrium model with Maxwellian free electrons. We run 10,000 PrismSPECT simulations, for electron temperatures linearly distributed in the range $T_e \subseteq [0.5, 25] \text{ eV}$, and ion density logarithmically distributed between $n_i \subseteq [1 \times 10^{16}, 1 \times 10^{19}] \text{ cm}^{-3}$.

Our in-house radiation transport solver computes the output intensity spectrum $I_\omega(s)$ given spatially-varying emissivity $\epsilon_\omega(s)$ and opacity $\alpha_\omega(s)$ values, by solving the steady-state radiation transport equation along the 1D path s [13]:

$$\partial_s I_\omega(s) = \epsilon_\omega(s) - \alpha_\omega(s) I_\omega(s) \quad (1)$$

Emissivity and opacity values vary along the LOS due to spatial variations in density and temperature. PrismSPECT




Fig. 2: Comparison of emissivity predictions made by a KNN model trained on PrismSPECT simulation results. We compare results for randomly-chosen members of the test set. Emissivity is scaled to $[0, 1]$.

computes ϵ_ω and α_ω for spatially-homogeneous plasmas. To construct the spatially-varying emissivity and opacity, we first assume some m -dimensional density and temperature distributions $n_i(s), T_e(s) \in \mathcal{R}^m$ along the LOS. Here, m is the number of points required to capture the spatial variation along the LOS. We then calculate the emissivity and opacity values at each position s along the LOS from the values of $n_i(s)$ and $T_e(s)$ at that location. The radiation transport solver then determines the output intensity distribution by solving Equation 1. We use a spectral resolution of 3.75×10^{-3} eV for our radiation transport calculations, which results in a $[400 \times 1]$ dimensional intensity spectrum for each data point.

In Figure 1, we compare the output intensity generated by the radiation transport solver with that from a planar plasma simulation in PrismSPECT. Here, the ion density and electron temperature of the plasma are $n_i = 5 \times 10^{17} \text{ cm}^{-3}$ and $T_e = 2.5 \text{ eV}$ respectively, while the size of the plasma is 10 mm. The planar plasma simulation includes the effect of radiation transport, for the case of a homogenous (i.e. constant density and temperature) slab of a specified size. The output of the radiation transport solver agrees with that from PrismSPECT. Figure 1 also shows the case for which the optical thickness is set to zero. For the $\tau = 0$ case, lines with high opacities are no longer damped by absorption from the plasma, and thus, the line ratios are significantly modified when compared to the case with radiation transport. This illustrates that in our pulsed-power-driven plasma of interest, optical thickness is important, and must be included in the spectroscopy analysis.




Fig. 3: (a) Simulated ion density at peak current, generated by a 40 mm diameter exploding wire array with 150 aluminum wires, driven by a 10 MA current pulse (300 ns rise time). This simulation was performed using GORGON, a two-temperature resistive MHD code (b) Variation of density and temperature along a chordal line-of-sight (LOS) as shown in (a).

The radiation transport solver requires emissivity and opacity calculated at each density and temperature value along the LOS as inputs to generate the intensity spectrum. We interpolate the emissivities and opacities for the intermediate values not simulated with PrismSPECT, using a k-nearest neighbor (KNN) regressor, trained on the output of the 10,000 PrismSPECT simulations. To evaluate the performance of the regressor, we compare the predicted spectra with the previously unseen emissivity spectra in the test set. As observed in Figure 2, where we compare the predicted emissivity with the actual emissivity for randomly-chosen members of the test set, the predictions agree well with the actual spectra. The coefficient of determination (also called the R2-score) is a commonly-used metric to characterize the performance of regression, and is defined as:

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2} \quad (2)$$

Here, y_i is the predicted value, $\hat{y}_i$ is the actual value, and $\bar{y} \equiv 1/n \sum_{i=0}^n y_i$ is the mean of the actual values. For the problem above, the KNN regressor exhibits an R2-score of 99.62%, showing that the model accurately reproduces the emissivity and opacity spectra for density and temperature values not included on the simulation grid in PrismSPECT.

B. Density and Temperature in Exploding Wire Arrays

Exploding wire arrays, which consist of a cylindrical cage of wires around a central cathode, are commonly used sources of pulsed-power-driven plasma for laboratory astrophysics experiments [14], [16], [17]. The magnetic field is oriented in the azimuthal direction inside the wires, and results in a $\mathbf{j} \times \mathbf{B}$ force that accelerates the ablating plasma radially outwards from the wires. Due to radially-diverging flows, the density decays rapidly with distance from the wires. This can be observed in Figure 3a, which shows the simulated ion density distribution generated by a 40 mm diameter exploding wire array with 150 aluminum wires, driven by a 10 MA current pulse with a 300 ns rise time. This simulation was performed using GORGON — a two-temperature Eulerian resistive magnetohydrodynamic (MHD) code [25].

As discussed before, we require m -dimensional arrays to fully capture the spatial variation in density and temperature along the LOS. However, we can use simplifying assumptions to make the problem more computationally tractable. If the emission is recorded along a chordal LOS, as shown in Figure 3a, the density variation can be approximated as Gaussian (see Figure 3b) i.e. $n_i(s) = n_0 \exp[(s - s_0)/2\sigma^2]$. Here, n_0 is the peak density, σ is the standard deviation of the Gaussian function, and we set the mean s_0 to half the total path length. The density distribution shown in Figure 3 is consistent that measured experimentally with laser imaging interferometry in previous pulsed-power experiments [16]. Our MHD simulations (Figure 3b) and previous experimental measurements also show little spatial variation in the temperature due to short thermal diffusion time in pulsed-power plasmas [17]. Therefore, we approximate the temperature to be constant along the LOS i.e. $T_e \neq T_e(s) = T_0$. This allows us to reduce




Fig. 4: Normalized spectral intensity simulated by the radiation transport solver. In (a)-(c), the blue curves correspond to output spectra generated for a Gaussian density variation ($n_0 = 5 \times 10^{17} \text{ cm}^{-3}$, $\sigma = 5 \text{ mm}$), and constant temperature ($T_e = 2 \text{ eV}$) along the spectroscopy LOS. (a) Change in the intensity spectrum with increasing temperature. (b) Change in the intensity spectrum with increasing density n_0 . (c) Change in the intensity spectrum with increasing σ .

our $2 \times m$ dimensional problem to just 3 variables — n_0 , T_0 , and σ . Both the density and temperature in Figure 3b also exhibit small amplitude modulations, which arise due to oblique shocks resulting from the azimuthal expansion of plasma from the discrete wires [26]. Our radiation transport calculations, however, show that the effect of these modulations on the recorded intensity spectrum is small.

Figure 4 shows a synthetic intensity spectrum, generated by the radiation transport solver, with values $n_0 = 5 \times 10^{17} \text{ cm}^{-3}$, $T_e = 2 \text{ eV}$, and $\sigma = 5 \text{ mm}$. Here, we normalize the spectrum between $[0, 1]$ by dividing by the maximum intensity. The spectrum exhibits Al-II and Al-III lines, which correspond to transitions in singly- (Mg-like) and doubly-ionized (Na-like) aluminum respectively. When the temperature is increased (Figure 4a), the relative intensity of the Al-III lines compared to the Al-II lines increases. This is expected because the ionization is higher at a higher temperature, and thus, the relative population of the higher- $\bar{Z}$ Al-III ions increases relative to the singly-ionized Al-II ions. The Al-II and Al-III lines only appear simultaneously between 1.5-3.5 eV. In Figure 4a, at 4 eV, the Al-II lines are completely suppressed. When we increase the density (Figure 4b), the lines not only become broader (due to Stark broadening) but the line ratios change as well. This is because increasing the density also increases the optical thickness, and the optically-thick lines are damped more strongly. This can be observed in Figure 4b when we compare the relative intensity of the Al-II 466.4 nm line (which is relatively optically thin) with that of the higher-

opacity Al-II 624.0 nm line. The more optically-thick Al-II 624.0 nm line is strongly damped at higher densities. Finally, changing the value of σ (Figure 4c), also changes the line ratios because the optical thickness increases with the size of the plasma; however, the sensitivity of the spectrum to changes in σ is relatively smaller than that in density and temperature.

To generate our training dataset for the machine learning task, we randomly sample values of $n_0 \subseteq [0.5 \times 10^{17}, 1 \times 10^{19}] \text{ cm}^{-3}$, $\sigma \subseteq [5, 40] \text{ mm}$, and $T_0 \subseteq [0.5, 25] \text{ eV}$ from a uniform distribution. Our radiation transport solver then uses these sampled values to calculate intensity spectra for each n_0 , σ , and T_e . We generate a total of 2500 $[400 \times 1]$ intensity spectra to use as training data for our 3-variable regression problem. Lastly, we scale and normalize the values of n_0 , σ and T_0 so that they lie within the interval $\subseteq [0, 1]$. We use linear scaling for the temperature and σ , and logarithmic scaling for the density. We also scale the intensity output of the radiation transport solver to the range $\subseteq [0, 1]$. This means we only use the shape of the intensity spectrum for our prediction, which obviates the need for absolute intensity calibration.

III. METHODOLOGY

A. Single-Objective Optimization

Our goal is to predict density and temperature profiles given a measured intensity spectrum I_{target} . One way to frame this problem is as a single-objective optimization problem:

$$\begin{aligned} \min_{\mathbf{x}} : f(\mathbf{x}) &= \text{MSD}(I_{\text{target}} - I(\mathbf{x})) \\ \text{where: } \mathbf{x} &= [n_0, \sigma, T_0] \\ \text{s.t. : } 0.5 \times 10^{17} &\leq n_0 [\text{cm}^{-3}] \leq 1 \times 10^{19} \\ 5 &\leq \sigma [\text{mm}] \leq 40 \\ 0.5 &\leq T_0 [\text{eV}] \leq 25 \end{aligned} \quad (3)$$

We minimize the mean squared deviation (MSD) between the target and predicted intensities. The objective function can be represented as:

$$\text{MSD} = \frac{1}{N} \sum_i^N [I_{\text{pred}}(\omega_i) - I_{\text{target}}(\omega_i)]^2 \quad (4)$$

Here, $I_{\text{pred}}(\omega_i)$ and $I_{\text{target}}(\omega_i)$ are the simulated and target intensities at frequency ω , and $N = 400$ is the size of the intensity spectrum. We use the radiation transport solver described in subsection II-A to generate the predicted intensities.

We perform the optimization using a $(\mu + \lambda)$ genetic algorithm (GA) implemented using the pymoo package in Python. The GA optimization algorithm iteratively searches for solutions that minimize the objective function over multiple generations [27]. In each generation, the best-performing solutions are selected and included in the population for the next iteration. Solutions are combined in each iteration in a process called crossover to create offspring solutions. The solutions are also subject to random changes in the values of the variables $\mathbf{x}$ to increase the diversity of the solutions, in a process called mutation. For the optimization here, we use a randomly generated initial population size of 150, with simulated binary crossover and polynomial mutation

(probability = 0.5, distribution index $\eta_c = 1$). We terminate the optimization when the MSD of the solutions becomes lower than a specified threshold. We repeat the optimization 50 times with different starting seeds to construct a family of optimal solutions. We exclude solutions from runs in which the GA gets stuck in a local minimum, where the objective function does not converge to a value below the required threshold.

B. Multi-Variable Regression

An alternative approach is to formulate the problem as a multi-variable regression problem. Given an input of an unseen intensity spectrum I_{target} , we predict the corresponding (normalized) values of n_0^* , T_0^* , and σ^* using machine learning-based regression models. Here, we compare the performance of multiple regressors for our three-variable regression task — linear regression (LR), k-nearest neighbor regressor (KNN), decision trees (DT), random forest (RF), deep neural network (DNN), and a 1D convolution neural network (1D-CNN). The choice of regression algorithm often represents a trade-off between model precision and interpretability. Simpler models such as linear regression and KNN models are relatively easier to understand and interpret, whereas deep neural networks and AutoML, which often provide high performance, require more training time, and are challenging to interpret [28].

The linear regression, KNN, decision trees, gaussian process regressor, and random forest models are implemented using the scikit-learn [29] package in Python. The KNN algorithm predicts values based on the distance from the k nearest data points in the training set [28]. Here, we use $k = 8$ and Euclidian distance for our KNN regressor. Decision trees follow a flowchart-like structure, and make predictions by asking questions at each level [28]. We use a DT regressor with a depth of 5 and a minimum sample split of 5. Random forests are ensembles of many decision trees [28]. Our random forest regressor uses 140 estimators; minimum samples required for a split is 4, and the minimum samples per leaf are 5.

We use the TensorFlow [30] package to construct a 3-layer deep neural network (DNN). Neural networks consist of multiple ‘hidden’ layers, consisting of several neurons, sandwiched between an input and an output layer [31]. In our DNN architecture, each layer is a fully-connected dense layer with 100 neurons (i.e. each neuron is connected to every other neuron in the previous and next layers), with ReLu activation functions, and padded with a batch normalization layer and a dropout layer ($p = 0.3$) to help prevent overfitting. The final layer consists of a 3-dimensional dense layer. Similarly, the 1D-CNN consists of two 1D convolution layers (kernel size = 3) with filter sizes of 8 and 16 respectively, and leaky ReLU activation layers. The convolution layer performs a convolution operation using the specified kernel on the input from the previous layer [31]. Each convolution layer is padded with a batch normalization layer, and a $p = 0.1$ dropout layer. The convolution layers are followed by a 200-neuron dense fully-connected layer, and a 3-dimensional output layer. Both neural networks use the Adam optimizer with a 0.5×10^{-4} learning rate, and a mean squared error (MSE) loss function. The key difference between the DNN and CNN architectures is that

the DNN treats the input vector as a 400-dimensional vector of parameters, whereas the CNN treats it as a 1D image, and therefore, has information about the relative spectral location of each element in the input vector.

Finally, we also implement a tabular AutoGluon model using Python’s AutoGluon package [32]. AutoGluon provides an automated approach to machine learning, by automatically comparing and combining the performance of many different models. The performance of AutoGluon has previously been shown to exceed that of more traditional ML models [32].

Each regressor is trained on the 2500×400 synthetic spectra with a 2000:500 split between the training and test sets, and $k = 3$ stratified k-fold cross-validation for the training set. Members of the test set are not shown to the ML model during the regression task, and are used to evaluate the regression performance after training. The hyperparameters described in this section were determined using hyperparameter optimization implemented using the Python package Optuna [33].

We use two metrics to characterize the performance of the different regression models. These are the coefficient of determination (also called the R2-score), and the mean squared deviation (MSD) from the simulated curve. The R2-score is defined in Equation 2, and is calculated from the predicted and actual values of n_0^* , T_0^* and σ^* . Similarly, the MSD (Equation 4) measures the deviation of the predicted intensity spectrum from the actual spectrum. Here, the predicted spectrum is determined by feeding the predicted values of n_0^* , T_0^* and σ^* from the ML model into the radiation transport solver. The R2-score and MSD are calculated for the $N = 500$ test set, which contains spectra not previously seen by our ML models. For good performance, we aim to




Fig. 5: (a) Comparison of target intensity spectra with spectra determined via optimization for the test case $n_0 = 1 \times 10^{18} \text{ cm}^{-3}$, $T_0 = 2 \text{ eV}$, $\sigma = 10 \text{ mm}$. The target spectrum is the solid blue curve, while the predicted spectrum and the predicted values are in orange. (b) Robustness of the fitting to noise and stray lines.

TABLE I: Comparison of the performance of different ML models on the regression task

Model	Training Time (s)	R2-Score(%)	2 Var. R2-Score(%)	Median MSD [$\times 10^{-4}$]
Linear Regression (LR)	0.3	-41.54 ± 80.76	64.00 ± 10.2	1.1
Decision Tree (DT)	0.35	53.22 ± 4.5	89.9 ± 2.57	5.3
k-Nearest Neighbor	0.02	67.89 ± 2.8	96.25 ± 1.8	0.4
Random Forest (RF)	6.41	67.06 ± 1.93	96.58 ± 0.5	0.4
Deep Neural Net. (DNN)	240.12	71.54 ± 1.2	94.52 ± 1.1	4.2
1D Conv. Neural Net. (1D-CNN)	260.17	65.34 ± 3.2	91.69 ± 1.4	5.7
AutoGluon	4760.11	74.20 ± 2.1	98.98 ± 1.4	0.3

maximize the R2-scores, and minimize the MSD. The R2-score may provide spurious performance metrics in case of non-unique solutions, whereas, the MSD characterizes how close the predicted spectra are to the actual spectra, allowing us to overcome this issue.

IV. RESULTS AND DISCUSSION

A. Optimization results

Using GA-based optimization, we predict values of n_0 , T_0 and σ for several test cases. Figure 5a compares the target intensity spectra with spectra determined via optimization for a randomly-selected test case ($n_0 = 1 \times 10^{18} \text{ cm}^{-3}$, $T_0 = 2 \text{ eV}$, $\sigma = 10 \text{ mm}$). The predicted values of n_0 , T_0 and σ reproduce the target spectrum well. For the test case shown in Figure 5a, the MSD from the target spectrum is roughly 6×10^{-5} , and the predicted values are $n_0 = 1.0 \times 10^{18} \text{ cm}^{-3} \pm 7\%$, $T_0 = 2.0 \text{ eV} \pm 1\%$, and $\sigma = 9.9 \text{ mm} \pm 20\%$. The range in the value of σ for the family of optimal solutions is comparatively higher than that in the density and temperature. This is consistent with our observation in Figure 4c, which shows that changes in σ generate linear changes in the optical depth, whereas those in n_0 and T_0 result in larger non-linear changes in opacity and relative intensities.

In many real situations, the experimental data that we want to fit may be noisier than the synthetic spectrum shown in Figure 5a. Furthermore, the spectrum may also be contaminated by stray lines, generated by impurities in the plasma, or by radiation emitted from other photoionized surfaces. In order to test the robustness of the prediction, we add noise and stray lines to the synthetic target spectrum (see Figure 5b). The predicted solution reproduces the target spectrum well, despite the added noise and stray radiation. In this case, the MSD of the optimum solutions is about 3×10^{-3} , which as expected, is larger than that for the smooth case. The predicted solutions are $n_0 = 1.2 \times 10^{18} \text{ cm}^{-3} \pm 20\%$, $T_0 = 2.0 \text{ eV} \pm 2\%$, and $\sigma = 9.9 \text{ mm} \pm 30\%$. There is a larger uncertainty in the solutions for the noisy target spectrum when compared to the smooth target spectrum (Figure 5a); however, the predicted solutions still include the values used to generate the test case.

Although optimization using GA provides good results, the computational time is high ($> 10 \text{ min}$ per prediction). This is primarily because the GA has to simulate a large number of potential candidates iteratively using our radiation transport model over several generations before convergence is achieved. Although we use GA as an example here, other optimization-based curve-fitting algorithms can also exhibit high computation times. This method can be highly effective with small data sets; however, it can be less attractive in cases




Fig. 6: Violin plots of the MSD of the predicted intensity spectra from the actual intensity spectra in the test set. The red lines represent the medians of the distribution, the blue rectangle represents the inter-quartile range, and the end-caps represent $1.5 \times$ the inter-quartile range. The blue-shaded regions show the shape of the full distribution.

where we must analyze large datasets of spectra, or when we require quick or real-time analysis of spectral information. The use of ML-based regression models, discussed in the next section, can be more useful in such situations.

B. Multi-variable regression using machine learning

The training of ML models can be computationally time-intensive; however, once the training is complete, large datasets can be evaluated rapidly. The training time of models typically depends on the complexity of the model. Table I compares the performance of the ML regressors described in subsection III-B. Here, the R2-score is computed using the predicted and actual values of all three variables, while we compute the two-variable R2-score using only the value of density n_0^* and temperature T_0^* . AutoGluon was the best-performing model, with an R2-score of $74.20\% \pm 2.1\%$. Since AutoGluon trains and compares the performance of multiple different models simultaneously, the training time was high compared to the other models. The next best-performing models (deep neural network, Random Forest, and k-nearest neighbor) exhibited R2-scores between $67 - 71\%$ and shorter training times. The two-variable R2 scores, calculated for density n_0^* and temperature T_0^* only, are roughly 30% higher than the three-variable R2-scores, which shows that the models predict n_0^* and T_0^* with better accuracy than the third variable σ^* . There is a larger uncertainty in the prediction of σ^* , because as observed in Figure 4c and in subsection IV-A, changes in σ^* generate linear changes in the optical depth,




Fig. 7: Distribution of predicted values of n_0^* , T_0^* and σ^* against the actual values for the AutoGluon Model. Here, the values of peak density, temperature, and the spread of the density function σ have been normalized. Blue points compare the values for the training set, while orange points do so for the test set.

whereas those in n_0^* result in larger non-linear changes. We also find that in cases where σ^* is under-predicted, the density is slightly over-predicted and vice versa, indicating that the model relies on small changes in density rather than on σ^* to incorporate the effects of changing optical depth. The majority of models, with the exception of linear regression and decision trees, exhibit two-variable R2-scores $> 90\%$, showing the effectiveness of these models in predicting n_0^* and T_0^* , despite the relatively poorer prediction of σ^* . AutoGluon, once again, exhibits the highest two-variable R2-score ($98.98\% \pm 1.4\%$).

In these models, we only use the shape of the intensity spectrum to make predictions. This can be advantageous as it obviates the need for absolute intensity calibration of the intensity spectra. However, adding information about the absolute intensity can potentially improve the predictions, as the absolute intensity also depends on the temperature, density, and size of the plasma. This would, however, require accurate modeling of intensity attenuation and losses in the optics used for the experiment. When we include the absolute intensity in the training data (that is, we do not normalize the intensity between $[0, 1]$), the models show improved prediction of σ , and the R2-scores increase by roughly 3-7%.

Figure 6 compares the MSD of the predicted intensity spectra from the actual intensity spectra in the test set for the different ML models. Here, we use the predictions of n_0 , T_0 and σ from the ML regressors as inputs into the radiation transport solver to determine the predicted spectra. The red lines represent the medians of the distributions, the blue rectangles represent the interquartile range, and the shaded regions show the shape of the full distribution. As expected, the predictions of the AutoGluon model, which exhibit the highest R2-score, also exhibit the lowest MSD, i.e. the predictions match the original intensity spectra well. The median MSD is relatively small ($< 0.1\%$) for all the models, however, for models with lower R2-scores, the distribution of the MSD exhibits a larger spread and extends to higher values. Interestingly, although the DNN exhibits R2-scores similar to the RF and KNN models, its MSD distribution is significantly wider. For the DNN, the distribution is also wider than for the 1D-CNN model, which exhibited lower R2-scores. This

may indicate that although the predictions of n_0 , T_0 and σ made by the DNN are close to the actual values, the values are consistently off, resulting in a relatively larger deviation between the predicted intensity calculated from these values and the actual intensity used to predict the values.

To gain further insight into the performance of the models, we plot the distribution of predicted values against the actual values. For the best-performing AutoGluon model (Figure 7), we find that predicted values of n_0^* and T_0^* from the test set exhibit the least deviation from the diagonal, whereas predicted values of σ^* exhibit a relatively larger deviation from the diagonal, as expected. However, AutoGluon is approximately able to capture the relative trend in the values of σ^* , as seen in the positive slope of the distribution in Figure 7b. This is in contrast to other models, which exhibit larger deviations in the predictions of σ^* . Figure 7a shows that the predicted values of density deviate from the actual values more significantly at very low ($n_0 < 1 \times 10^{17} \text{ cm}^{-3}$) and very high densities ($n_0 > 5 \times 10^{18} \text{ cm}^{-3}$). At low density, the effects of optical depth and Stark broadening are smaller, which leads to larger uncertainties in the prediction of the density. Our radiation transport simulations also show that at higher densities, line radiation is strongly damped due to the opacity effects, and continuum emission begins to dominate. This may explain why the predicted density deviates more significantly at these higher densities. For the temperature predictions (Figure 7c), the deviations are relatively small for $T_0 < 3 \text{ eV}$, and become more significant at higher temperatures. Below 3 eV, as seen in Figure 4a, the spectrum contains both Al-II and Al-III lines, and the relative intensities of the inter-stage lines are a strong function of temperature. However, at higher temperatures, only the Al-III lines are present in the spectrum, and continuum emission also begins to dominate, which makes it relatively harder to predict T_0 .

A key challenge with more complicated ML models, such as neural networks and AutoGluon, is the lack of interpretability. To gain insight into features that inform the prediction, we investigate the relative importance of different parts of the intensity spectrum for the prediction of density and temperature. In order to do so, we calculate the sensitivity of the

density and temperature predictions to perturbations in the spectral intensity value at different wavelengths. Figure 8 shows the sensitivity map computed for the AutoGluon model for a randomly-selected member of the test set. For this intensity spectrum, the actual values of density, temperature, and σ are $1.4 \times 10^{18} \text{ cm}^{-3}$, 7.3 eV, and 25 mm, and the predicted values are $1.5 \times 10^{18} \text{ cm}^{-3}$, 6.9 eV, and 22 mm. We approximate the gradients in density and temperature ($\partial n_0^*/\partial \epsilon$ and $\partial T_0^*/\partial \epsilon$) by perturbing the input intensity spectrum at a given wavelength by a small value ϵ , and dividing the change in the predicted value ($|\delta n_0^*|$ or $|\delta T_0^*|$) by the perturbation ϵ . Here, we pick the perturbation ϵ from a Gaussian distribution of amplitude 0.1, and the mean gradients are calculated over multiple iterations to determine the final value. Such sensitivity maps are commonly used in image classification problems to identify parts of an image that may contribute to the final classification [34]. Here, we use it to identify features of the spectra that contribute to the final predictions in our regression problem. In Figure 8, as expected, the lines (Al-III 415.2 nm, Al-III 448.1 nm, Al-III 452.5 nm, and Al-III 570.7 nm) contribute significantly to the prediction of density and temperature, whereas parts of the spectrum that correspond to the continuum are less important. This can be observed from the large gradients $\partial n_0^*/\partial \epsilon$ and $\partial T_0^*/\partial \epsilon$ at the positions of these lines. The Al-III 570.7 nm line appears to be particularly important for the temperature prediction, while the Al-III 448.1 nm and Al-III 452.5 nm lines (which have merged at this higher density) are relatively more important to the density prediction. In addition to the locations of the Al-III lines, smaller peaks in the gradients, particularly in temperature, also appear at parts of the spectrum devoid of lines. The locations of these peaks correspond to Al-II emission lines that appear predominantly at lower temperatures (0.5–1.5 eV), indicating that the absence of these lines contributes, although less significantly, to the temperature prediction.

Another limitation of these models is the reliance on synthetic data, which, in turn, is affected by the uncertainty and assumptions in the theoretical modeling used to generate the synthetic dataset. Benchmarking the results using independent diagnostics can be one way to probe the applicability of these ML models to real experimental data. However, as




Fig. 8: Sensitivity of the predictions of density and temperature to variations in the spectral intensity value at different wavelengths for a randomly-chosen intensity spectrum from the test set.

mentioned earlier, independent measurements are not always possible, and the uncertainty in CR modeling is inherent in the analysis of most spectroscopic data. As discussed in subsection IV-A, the experimental data can also be noisy, and include background radiation and contamination by impurities. When we use the trained ML models to make predictions from noisy spectra contaminated with stray radiation, the R2-scores typically fall by 10 – 20%. For a given experiment, the instrument response, the bit depth of the spectrometer, and attenuation by the optics may also need to be properly included in the synthetic dataset. The analysis of experimental spectra using these methods will be pursued in a future publication.

V. CONCLUSIONS

We explore the use of machine learning (ML) methods for rapid spectroscopic analysis of emission spectra in the visible regime. Our goal is to predict density and temperature in a pulsed-power-driven aluminum plasma generated by an exploding wire array. In contrast to previous work, which has typically focused on low-density homogenous plasmas, we aim to diagnose a high energy density non-optically thin plasma, which necessitates the use of radiation transport calculations to accurately model the recorded intensity spectrum. These radiation transport calculations use spectral emissivity and opacity values computed using the collisional-radiative code PrismSPECT. Consistent with previous observations in exploding wire arrays, we assume a constant temperature and a Gaussian density variation (peak density n_0 and standard deviation σ) along the diagnostic line of sight.

The radiation transport solver is first used to directly solve a single-objective optimization problem, which varies the values of n_0 , T_0 , and σ to minimize the mean squared deviation between a generated spectrum and the target spectrum. This approach provides reliable fits to the target spectrum, robust to noise and stray line radiation; however, it can be time-intensive, which limits its usefulness for rapid or real-time analysis of large datasets.

We then use the radiation transport solver to generate a $[2500 \times 400]$ dataset of synthetic emission spectra, and we compare the performance of different ML models on their ability to predict density, temperature, and σ from the given intensity spectra. The AutoGluon model performs best, with an R2-score of roughly 98% for density and temperature predictions. Simpler models (random forest, k-nearest neighbor, and deep neural network) also exhibit high R2-scores ($> 90\%$) for density and temperature predictions, showing the potential of ML models in providing rapid and accurate analysis of spectral data. However, the prediction of σ is relatively poor, which is typical of radiation transport problems, and is related to the relatively smaller sensitivity of the optical depth on this parameter, when compared to the peak density n_0 .

The mean square deviation from the actual spectra of the ML predictions is typically larger than that of the more time-intensive optimization task. One way to improve the accuracy of these models could be to perform further optimization by including the ML predictions in the initial population; this can provide faster convergence of the optimization and better fitting to the target spectra.

ACKNOWLEDGMENT

This work was funded by NSF and NNSA under grant no. PHY2108050, by the EAGER grant no. PHY2213898, and supported by the MathWorks fellowship.

REFERENCES

- [1] H. R. Griem, *Principles of plasma spectroscopy*, 2005.
- [2] U. Fantz, “Basics of plasma spectroscopy,” *Plasma sources science and technology*, vol. 15, no. 4, p. S137, 2006.
- [3] X.-M. Zhu and Y.-K. Pu, “Optical emission spectroscopy in low-temperature plasmas containing argon and nitrogen: determination of the electron temperature and density by the line-ratio method,” *Journal of Physics D: Applied Physics*, vol. 43, no. 40, p. 403001, 2010.
- [4] I. H. Hutchinson, *Plasma Phys. Control. Fusion*, vol. 44, no. 12, p. 2603, 2002.
- [5] G. A. Rochau, J. Bailey, and Y. Maron, *Phys. Plasmas*, vol. 17, no. 5, p. 055501, 2010.
- [6] M. R. Gomez, S. B. Hansen, K. J. Peterson, D. Bliss, A. Carlson, D. Lamppa, D. Schroen, and G. Rochau, “Magnetic field measurements via visible spectroscopy on the z machine,” *Review of Scientific Instruments*, vol. 85, no. 11, 2014.
- [7] H. R. Griem, “Plasma spectroscopy in inertial confinement fusion and soft x-ray laser research,” *Physics of Fluids B: Plasma Physics*, vol. 4, no. 7, pp. 2346–2361, 1992.
- [8] S. Sudkewer, “Spectroscopic diagnostics of tokamak plasmas,” *Physica Scripta*, vol. 23, no. 2, p. 72, 1981.
- [9] D. Astic, A. Canosa, P. Vigliano, and M. Autric, “Visible spectroscopy of laser produced aluminum plasma. electron temperature and density determinations,” in *International Congress on Applications of Lasers & Electro-Optics*, vol. 1991, no. 2. Laser Institute of America, 1991, pp. 94–103.
- [10] J. T. Knudtson, W. B. Green, and D. G. Sutton, *Journal of Applied physics*, vol. 61, no. 10, pp. 4771–4780, 1987.
- [11] J. Bailey, G. Rochau, R. Mancini, C. Iglesias, J. MacFarlane, I. Golovkin, J. Pain, F. Gilleron, C. Blancard, P. Cosse *et al.*, “Diagnosis of x-ray heated mg/fe opacity research plasmas,” *Review of Scientific Instruments*, vol. 79, no. 11, 2008.
- [12] J. MacFarlane, I. Golovkin, P. Woodruff, S. Kulkarni, and I. Hall, “Simulation of plasma ionization and spectral properties with prismspect,” *2013 Abstracts IEEE International Conference on Plasma Science (ICOPS)*, pp. 1–1, 2013.
- [13] R. P. Drake, *Phys. Today*, vol. 63, no. 6, p. 28, 2010.
- [14] S. Lebedev, A. Frank, and D. Ryutov, *Rev. Mod. Phys.*, vol. 91, no. 2, p. 025002, 2019.
- [15] K. Bell, T. Shelkovenko, S. Pikuz, R. McBride, I. Blesener, P. Knapp, D. Hammer, J. Greenly, and Y. Maron, “Optical spectroscopy experiments on the 500 ka xp pulsed-power generator,” *AIP Conference Proceedings*, vol. 1088, no. 1, pp. 161–163, 2009.
- [16] R. Datta, D. R. Russell, I. Tang, T. Clayson, L. G. Suttle, J. P. Chittenden, S. V. Lebedev, and J. D. Hare, *Journal of Plasma Physics*, vol. 88, no. 6, p. 905880604, 2022.
- [17] D. R. Russell, G. C. Burdiak, J. J. Carroll-Nellenback, J. Halliday, J. Hare, S. Merlini, L. G. Suttle, V. Valenzuela-Villasaca, S. Eardley, J. A. Fullalove *et al.*, “Perpendicular subcritical shock structure in a collisional plasma experiment,” *Physical Review Letters*, vol. 129, no. 22, p. 225001, 2022.
- [18] R. Datta, J. Angel, J. Greenly, S. Bland, J. Chittenden, E. Lavine, W. Potter, D. Robinson, E. Wong, D. Hammer *et al.*, “Plasma flows during the ablation stage of an over-massed pulsed-power-driven exploding planar wire array,” *arXiv preprint*, 2023.
- [19] J.-H. Park, J.-H. Cho, J.-S. Yoon, and J.-H. Song, *Coatings*, vol. 11, no. 10, p. 1221, 2021.
- [20] T. S. Mansouri, H. Wang, D. Mariotti, and P. Maguire, *Journal of Physics D: Applied Physics*, vol. 55, no. 22, p. 225205, 2022.
- [21] K. Shojaei and L. Mangolini, *Journal of Physics D: Applied Physics*, vol. 54, no. 26, p. 265202, 2021.
- [22] D. Clayton, K. Tritz, D. Stutman, R. Bell, A. Diallo, B. LeBlanc, and M. Podestà, *Plasma Phys. Control. Fusion*, vol. 55, no. 9, p. 095015, 2013.
- [23] S. H. Lee, J. Lee, I. Yamada, and J. S. Park, *Rev. Sci. Instrum.*, vol. 87, no. 11, p. 11E533, 2016.
- [24] C. Samuelli, A. Mclean, C. Johnson, F. Glass, and A. Jaervinen, *Rev. Sci. Instrum.*, vol. 92, no. 4, p. 043520, 2021.
- [25] A. Ciardi, S. V. Lebedev, A. Frank, E. G. Blackman, J. P. Chittenden, C. Jennings, D. J. Ampleford, S. Bland, S. C. Bott, J. Rapley *et al.*, “The evolution of magnetic tower jets in the laboratory,” *Physics of Plasmas*, vol. 14, no. 5, 2007.
- [26] G. Swadling, S. Lebedev, N. Niasse, J. Chittenden, G. Hall, F. Suzuki-Vidal, G. Burdiak, A. Harvey-Thompson, S. Bland, P. De Grouchy *et al.*, “Oblique shock structures formed during the ablation phase of aluminium wire array z-pinches,” *Physics of Plasmas*, vol. 20, no. 2, 2013.
- [27] S. Mirjalili and S. Mirjalili, “Genetic algorithm,” *Evolutionary Algorithms and Neural Networks: Theory and Applications*, pp. 43–55, 2019.
- [28] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An introduction to statistical learning*. Springer, 2013, vol. 112.
- [29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, “Scikit-learn: Machine learning in python,” *Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.
- [30] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, “Tensorflow: a system for large-scale machine learning,” in *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, 2016, pp. 265–283.
- [31] J. A. Anderson, *An introduction to neural networks*. MIT press, 1995.
- [32] N. Erickson, J. Mueller, A. Shirkov, H. Zhang, P. Larroy, M. Li, and A. Smola, “Autogluon-tabular: Robust and accurate automl for structured data,” *arXiv preprint arXiv:2003.06505*, 2020.
- [33] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, “Optuna: A next-generation hyperparameter optimization framework,” in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 2623–2631.
- [34] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-cam: Visual explanations from deep networks via gradient-based localization,” *Proceedings of the IEEE international conference on computer vision*, pp. 618–626, 2017.