

# ESPNN: A novel electronic stopping power neural-network code built on the IAEA stopping power database. I. Atomic targets

Cite as: J. Appl. Phys. **132**, 245103 (2022); <https://doi.org/10.1063/5.0130875>

Submitted: 15 October 2022 • Accepted: 01 December 2022 • Published Online: 22 December 2022

 F. Bivort Haiek,  A. M. P. Mendez,  C. C. Montanari, et al.



View Online



Export Citation



CrossMark



## APL Quantum

CALL FOR APPLICANTS

Seeking Editor-in-Chief

# ESPNN: A novel electronic stopping power neural-network code built on the IAEA stopping power database. I. Atomic targets

Cite as: J. Appl. Phys. **132**, 245103 (2022); doi: [10.1063/5.0130875](https://doi.org/10.1063/5.0130875)

Submitted: 15 October 2022 · Accepted: 1 December 2022 ·

Published Online: 22 December 2022



F. Bivort Haiek,<sup>1</sup>  A. M. P. Mendez,<sup>2</sup>  C. C. Montanari,<sup>2</sup>  and D. M. Mitnik<sup>2,a)</sup> 

## AFFILIATIONS

<sup>1</sup>Minería de Datos y Descubrimiento del Conocimiento, Universidad de Buenos Aires, Buenos Aires, Argentina

<sup>2</sup>Instituto de Astronomía y Física del Espacio, CONICET, Universidad de Buenos Aires, Buenos Aires, Argentina

<sup>a)</sup>Author to whom correspondence should be addressed: [dmitnik@df.uba.ar](mailto:dmitnik@df.uba.ar)

## ABSTRACT

The International Atomic Energy Agency (IAEA) stopping power database is a highly valued public resource compiling most of the experimental measurements published over nearly a century. The database—accessible to the global scientific community—is continuously updated and has been extensively employed in theoretical and experimental research for more than 30 years. This work aims to employ machine learning algorithms on the 2021 IAEA database to predict accurate electronic stopping power cross sections for any ion and target combination in a wide range of incident energies. Unsupervised machine learning methods are applied to clean the database in an automated manner. These techniques purge the data by removing suspicious outliers and old isolated values. A large portion of the remaining data is used to train a deep neural network, while the rest is set aside, constituting the test set. The present work considers collisional systems only with atomic targets. The first version of the ESPNN (electronic stopping power neural-network code), openly available to users, is shown to yield predicted values in excellent agreement with the experimental results of the test set.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0130875>

## I. INTRODUCTION

At the end of 2015, the Nuclear Data Section of the International Atomic Energy Agency<sup>1</sup> (IAEA) inherited the monumental work done by Paul.<sup>2–7</sup> He collected about 1000 experimental stopping power measurements made in multiple laboratories worldwide, including publications from as early as 1928. He kick-started his database project in 1990 at the University of Linz, and it has been available to the scientific community since then. The IAEA assumed the responsibility of maintaining, updating, and disseminating this data collection,<sup>6</sup> which included tables, figures, and comparisons of the published stopping data for ions in atomic targets, compounds, and new materials of technological interest. An overview of the database contents can be found in the review of Montanari and Dimitriou.<sup>8</sup>

The stopping power is the mean energy loss per unit path length of the projectile in many collisional processes. Calculating the electronic stopping power involves determining the target system probabilities of occupying any electronic state different from

the initial one due to the transfer of energy from the ion to the target electrons. Several reviews on this subject are available in the literature.<sup>9–11</sup> Different methods and semiempirical codes freely available online are linked and scrutinized on the IAEA stopping power website.<sup>12</sup> Certainly, the code most widely used is SRIM. This code is based on a semiempirical method developed by Ziegler.<sup>13</sup> As reported in their work,<sup>14</sup> it reproduces 64% of the data with an overall accuracy of 5%. Noteworthy, the latest version of SRIM includes measurements only up until 2013. Essential differences between the SRIM predictions and new measurements have been reported since then.<sup>15–18</sup>

The present work is the first of a series of publications where we design a robust and general model to accurately calculate the electronic stopping power in different target materials and along an extended energy range. To fulfill this task, we developed a machine learning (ML) model based on a clustering technique and a deep neural-network (NN) method.

The IAEA database was built by gathering published articles from diverse authors; hence, it needs to be standardized in its

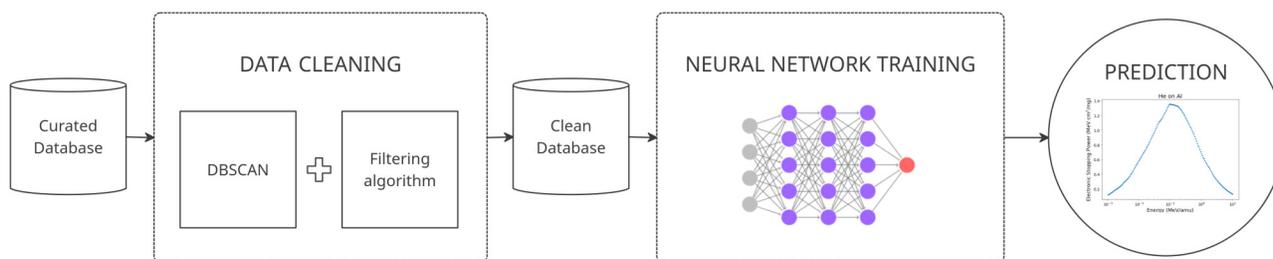


FIG. 1. Graphic scheme of the electronic stopping power with neural-network (ESPNN) model. See the text for details.

original form. The data are presented in various units and formats. Depending on the ion–target system, the experimental values are given as stopping power per unit length or cross sections per mass or atom. The database contains dozens of thousands of input values only for mono-elemental targets. As a preliminary work, we devoted significant efforts to reorganizing the database, unifying the units, and arranging the data in a standard (csv) format, which enabled easy and quick access to the compiled data. The first task consisted of cleaning the curated data outside the general trend. Purging these data by hand requires considerable amounts of work and is not recommended. Instead, we developed an unsupervised-machine-learning-based method to clean up the database by implementing a filtering algorithm and a cluster analysis. This clustering technique, called DBSCAN, identifies outlier values and determines which data to keep in the cases of inconsistent overlapping. The cleaning procedure is shown schematically in the left dashed-box of Fig. 1.

As illustrated in Fig. 1, the cleaned database becomes the input of the second supervised-machine-learning method consisting of a deep neural network. This network has many basic units (neurons) arranged in layers. Each neuron receives input signals from the previous elements, processes it through some weighted non-linear function, and transmits the resulting output to neurons belonging to the next layer. The neural network contains a known input (projectiles, targets, energies) and a known output (the corresponding experimental stopping power). The training is performed by adjusting the weights, minimizing the differences between the final processed output of the network (predictions) and the experimental results. In that way, the model can accurately reproduce the experimental values and, hopefully, predict new results in the cases not included in the training procedure (the test set). The resulting model and code ESPNN (electronic stopping power with neural network) are presented in this work. In this first article, we report the results obtained only for atomic targets. Our results were obtained with excellent accuracy using different error metrics, such as MAPE (Mean Average Percentage Error) or MAE (Mean Absolute Error).

Section II describes the machine learning (ML) methods employed to deplete the database, showing a few examples of the original data and the remaining final input employed in the network training. In Sec. III, the deep neural-network architecture is described. An overview of the training, validation, and test sets is presented; some details concerning the training procedure are also discussed. Section IV presents selected results and the analysis of

distinct error metrics, which give a sense of the method’s accuracy. In Appendix B, we provide instructions for installing and using the ESPNN code.

## II. DATABASE CLEANSING

### A. Database review

The database (updated in December 2021) consists of 60 173 experimental measurements, representing stopping power values for 1491 ion–target combinations of 49 projectiles colliding with 283 targets across the energy range  $10^{-4}$ – $10^4$  MeV/amu and ion and target atomic masses from 1 to 240 amu. Concerning only the mono-elemental targets, there are 706 collision cases composed of 44 different projectiles and 73 targets, resulting in 36 544 experimental data points. The experimental data summarize 1190 publications covering the period 1928–2021. The reorganization of the database allows for performing extensive statistical analysis. This task would indicate, for instance, the lack of data in certain energy regions, over-measured and under-measured systems, which may guide experimental groups regarding the necessity of new findings. A detailed analysis of the database’s current status will be presented in a forthcoming article.

The raw data collected come from several publications; the results of the same ion–target collision may show significant discrepancies (much larger than the error of the individual set of experimental data). Cleaning the database is crucial; well-thought criteria must be adopted for selecting the most reliable values, accompanied by careful examination of the outcome. The immense difficulties of scrutinizing such an extensive dataset are resolved by implementing a straightforward ML-based method. First, the DBSCAN classification algorithm is used to group similar results in clusters and to identify outliers, i.e., values suspected of being erroneous. Then, an algorithm is developed to assess clusters and outliers by introducing different criteria for overlapping and isolated data. In the following, we briefly explain these algorithms.

### B. The DBSCAN algorithm

Clustering algorithms are attractive for the task of class identification in spatial databases. However, most well-known unsupervised classification algorithms suffer severe drawbacks when applied to large spatial databases. That is, elements in the same cluster may not share enough similarities or performance may be

poor. Also, while partition-based algorithms, such as K-means, may be easy to understand and implement in practice, the algorithm has no notion of outliers; all points are assigned to a cluster, even if they do not belong to any. Moreover, anomalous points draw the cluster's centroid toward them, making it more difficult to classify them as anomalous points. In contrast, density-based clustering locates regions of high density that are separated by regions of low density (in this context, density is defined as the number of points within a specified radius).

In this work, we used the density-based spatial clustering of applications with noise algorithm DBSCAN.<sup>19,20</sup> The main idea behind this technique is the following: for a set of points in some space, the algorithm groups together values that are closely packed (points with many nearby neighbors), marking as outlier points that lie alone in low-density regions (whose nearest neighbors are too far away). It requires two input parameters: the radius of the neighborhood,  $\epsilon$ , and the number of reachable points (within a distance  $\epsilon$ ),  $N_{\min}$ , required to form a dense region. All points belonging to an  $\epsilon$ -neighborhood configure a cluster. All points not reachable from any other point are outliers or noise points. DBSCAN is significantly effective in discovering clusters of arbitrary shapes, which makes it a standard clustering algorithm and one of the most cited in the scientific literature. This method has numerous advantages: it does not require one to specify the number of clusters in the data *a priori* (as opposed to K-means, for example). DBSCAN can find arbitrarily shaped clusters; it has a notion of noise, is robust to outliers, and is mainly insensitive to the data order.

As examples of the clustering partitions produced by DBSCAN, we show in Fig. 2 the stopping power cross sections for

H in Si (left) and for O in Au (right). In all cases, the cross sections have been scaled to keep both axes of similar size. We found it convenient to set the input parameters as  $\epsilon = 0.025$  and  $N_{\min} = 3$ . For H in Si, the unsupervised algorithm detected 19 clusters (each one plotted with a different color). The small black dots without  $N_{\min}$  neighboring points inside a circle of radius  $\epsilon$  are considered noise points. In total, the algorithm labeled 139 outliers. For O in Au, 26 clusters and 181 outliers were identified.

### C. Filtering procedure

When the DBSCAN algorithm is used for data cleansing, the assigned outliers are generally and systematically removed from the dataset. However, the amount of data to be cleaned by following this standard course of action would be significant. Instead, we designed and implemented a three-step sequential filtering algorithm based on the clustering results.

Once the clusters and the outliers are identified, the sequential algorithm determines which data points are kept. First, we inspect energy regions with single and multiple publications based on a year-of-publication criterion. For example, in Fig. 2, the data points marked as outliers by DBSCAN may cover energy regions with no other available values. This context is essential and is considered by our cleaning algorithm: experimental values belonging to an energy region with no other measurements are considered valid data points; diversely, they are examined by the date criterion. In the following step, we examine overlapping (in energy) but isolated (in cross section) data points. Only those fulfilling an *ad hoc* criterion are kept; otherwise, they are tagged as true outliers. Finally, we inspect clustering composition, i.e., whether a cluster comprises

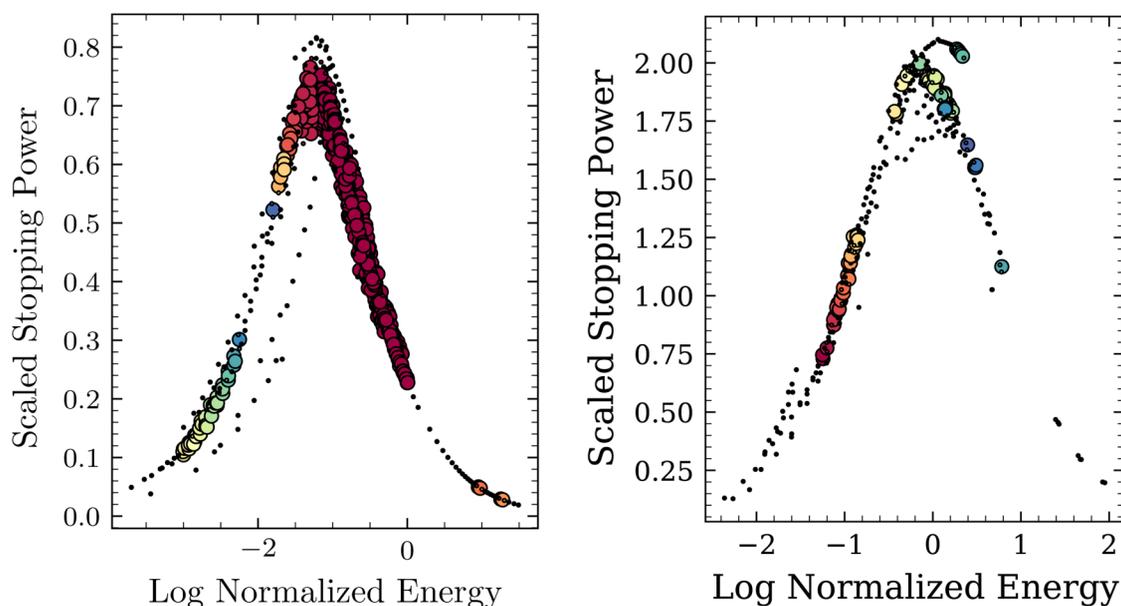


FIG. 2. Scaled stopping power cross sections for H in Si (left) and O in Au (right). The different colors represent clusters detected by the DBSCAN algorithm with  $\epsilon = 0.025$  and  $N_{\min} = 3$ . Outliers points are indicated with small dots.

data from single or many publications. The general structure of this cleaning procedure is described in the following, while a pseudo-code scheme resuming the filtering algorithm is presented in [Appendix A](#).

For each collisional system, every publication is first assessed according to the energy overlapping with other references. We determine the number of experimental values  $n_i$  published in the publication  $P_i$ , which are spread over an energy range  $\Delta E_i$ . If the publication data cover a region with no other measurements, these data are included in the input database. Otherwise, a date criterion is followed; for the remaining publications of the same collisional system, we count the experimental values that (1) fall inside the energy range  $\Delta E_i$  and (2) have been published after  $P_i$ . The total energy range covered by these posterior publications is denoted as  $\Delta E_i^p$ . We evaluate what portion of the energy range covered by the publication  $P_i$  is also covered by newer publications; i.e.,

$$\text{If } \frac{\Delta E_i^p}{\Delta E_i} \leq \sigma_\Delta \rightarrow \text{KEEP the publication } P_i. \quad (1)$$

The chosen overlap parameter  $\sigma_\Delta$  is close to 0.6 except for the cases in which the number of references addressing the specific collisional system is tiny; in such cases,  $\sigma_\Delta \approx 1$ . The condition in (1) ensures that if a publication presents results in an unexplored energy range, these values are kept and included in the input of the neural network, no matter if they belong to a cluster or are classified outliers. An example of this criterion is shown in the right panel of [Fig. 2](#); the experimental low energy values for O in Au are kept in the input database even though DBSCAN considers them as outliers.

Next, we deal with the remaining energy-overlapping results. For a given ion–target system, the second filtering step consists of dropping “isolated” results, i.e., data points appearing at the same energy range as others but having different values. In this instance, we count the number of outliers  $N_{\text{out}}^i$  detected in  $P_i$  and evaluate

$$\text{If } \frac{N_{\text{out}}^i}{n_i} > \sigma_{\text{out}} \rightarrow \text{DROP the publication } P_i, \quad (2)$$

where  $\sigma_{\text{out}} \approx 0.45$  except for when the number of publications is small; then,  $\sigma_{\text{out}} \approx 1$ . This criterion implies that if a publication has many results overlapping in energy with newer measurements, but most of their values are considered suspiciously wrong, it is convenient to keep the newer values rather than the outliers.

Finally, we perform the same test for isolated results, but instead of inspecting the outliers, we look into the clusters. We identify the largest cluster  $C_i$  in a particular publication  $P_i$ . This cluster has a total of  $t^i$  data points, including the  $l^i$  values belonging only to  $P_i$ . We aim to verify if the biggest cluster of a publication is formed by results obtained from many sources and not just from this unique work. Thus, we evaluate the condition

$$\text{If } \frac{l^i}{t^i} > \sigma_{\text{clu}} \rightarrow \text{DROP the publication } P_i, \quad (3)$$

with  $\sigma_{\text{clu}} \approx 0.45$  except for cases where the number of publications is small; in such cases,  $\sigma_{\text{clu}} \approx 1$ .

The performance of the cleaning procedure is illustrated with many examples in [Figs. 3–6](#). The left panels of all these figures display the complete data set for each collisional system. The central panels illustrate the cleaned data set, i.e., the outcome of implementing the cleaning procedure.

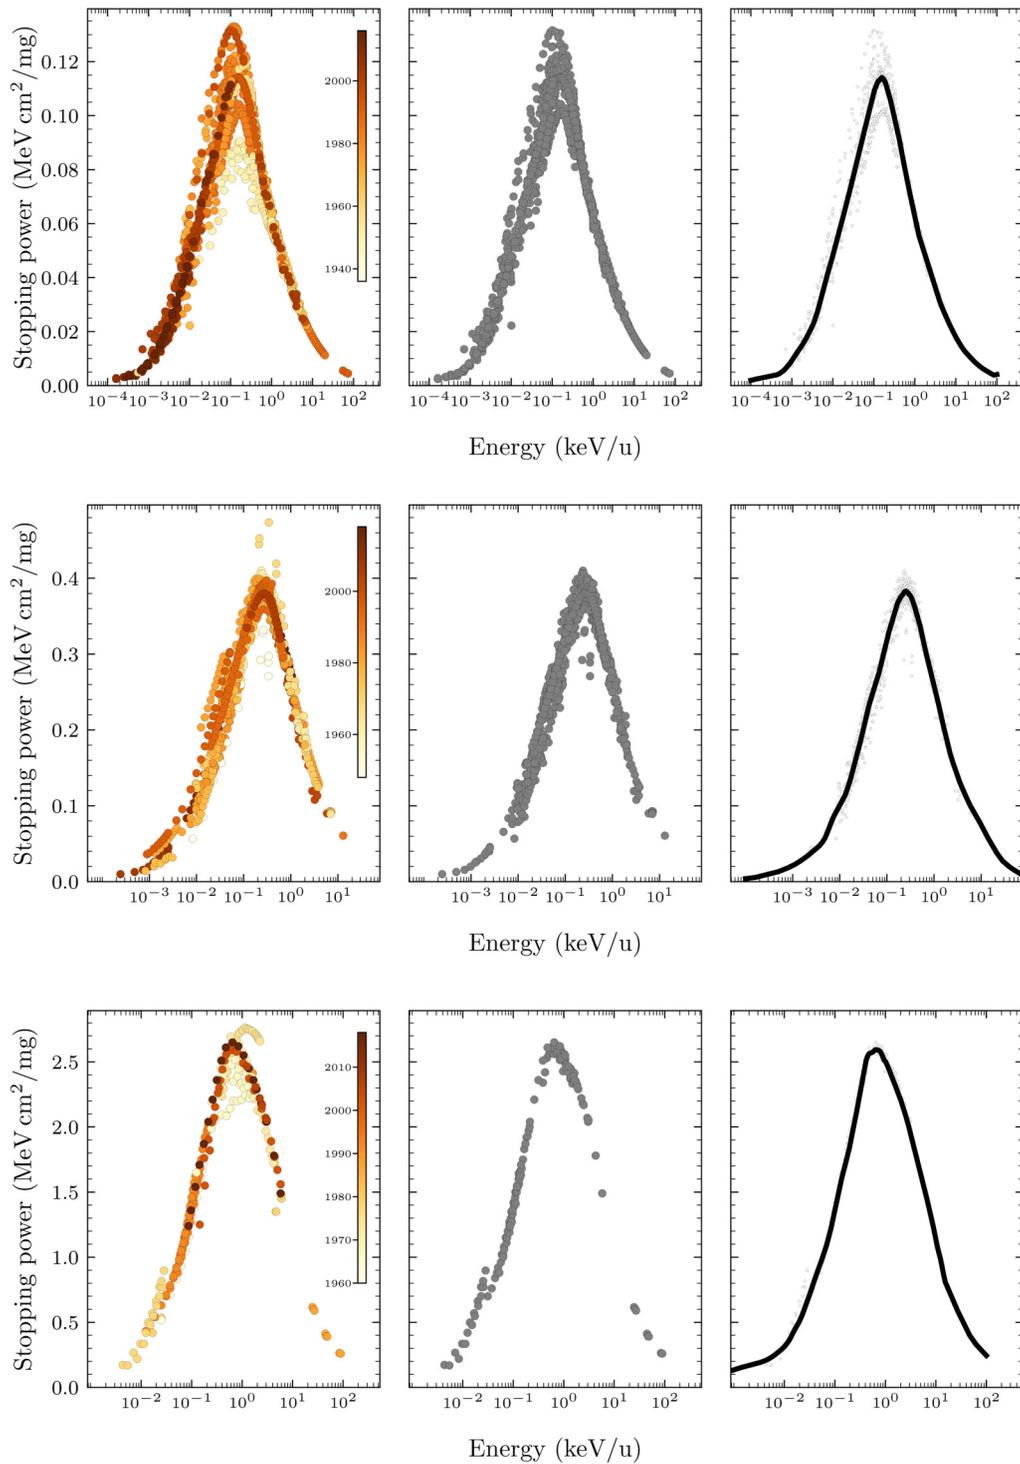
In [Fig. 4](#), we show the electronic stopping power cross section of projectiles H (top row), He (middle row), and O (bottom row), colliding with Au. For this target, a large amount of data is available. However, for these collisional systems, significant discrepancies among the various experimental results are evident, stressing the importance of an effective filtering mechanism. For H in Au, the top-center subplot shows that the cleaning algorithm can safely discard only the data produced before 1980. The middle-center subplot indicates that many results have been dropped for He in Au, particularly in the low energy range. The left subplot for O in Au illustrates multiple clusters overlapping in the energy region close to the stopping power cross section maxima. In the bottom-center subplot, we observe that the algorithm chooses to keep the latest experiments, discarding the old data; as a result, a clean and smooth curve is obtained.

It is important to stress that the filtering algorithm favors the latest data, but that does not necessarily imply removing the older results. The procedure will not rule out older measurements if they are statistically valid. Of course, the data cleansing procedure can be more severe or relaxed according to the parameters chosen. It is possible to change the DBSCAN parameters to modify the size of the clusters and the filtering parameters from Eqs. (1)–(3) for keeping or discarding some publications. Considering the success of the neural-network implementation that will be explored in the following, we assume that the values adopted are satisfactory for general purposes. An example of this feature is given in [Fig. 4](#), where we display the electronic stopping power cross sections of H colliding with Ni (top row), Cu (middle row), and Zn (bottom row). In these cases, the filtering algorithm does not remove the (most) differences near the maxima.

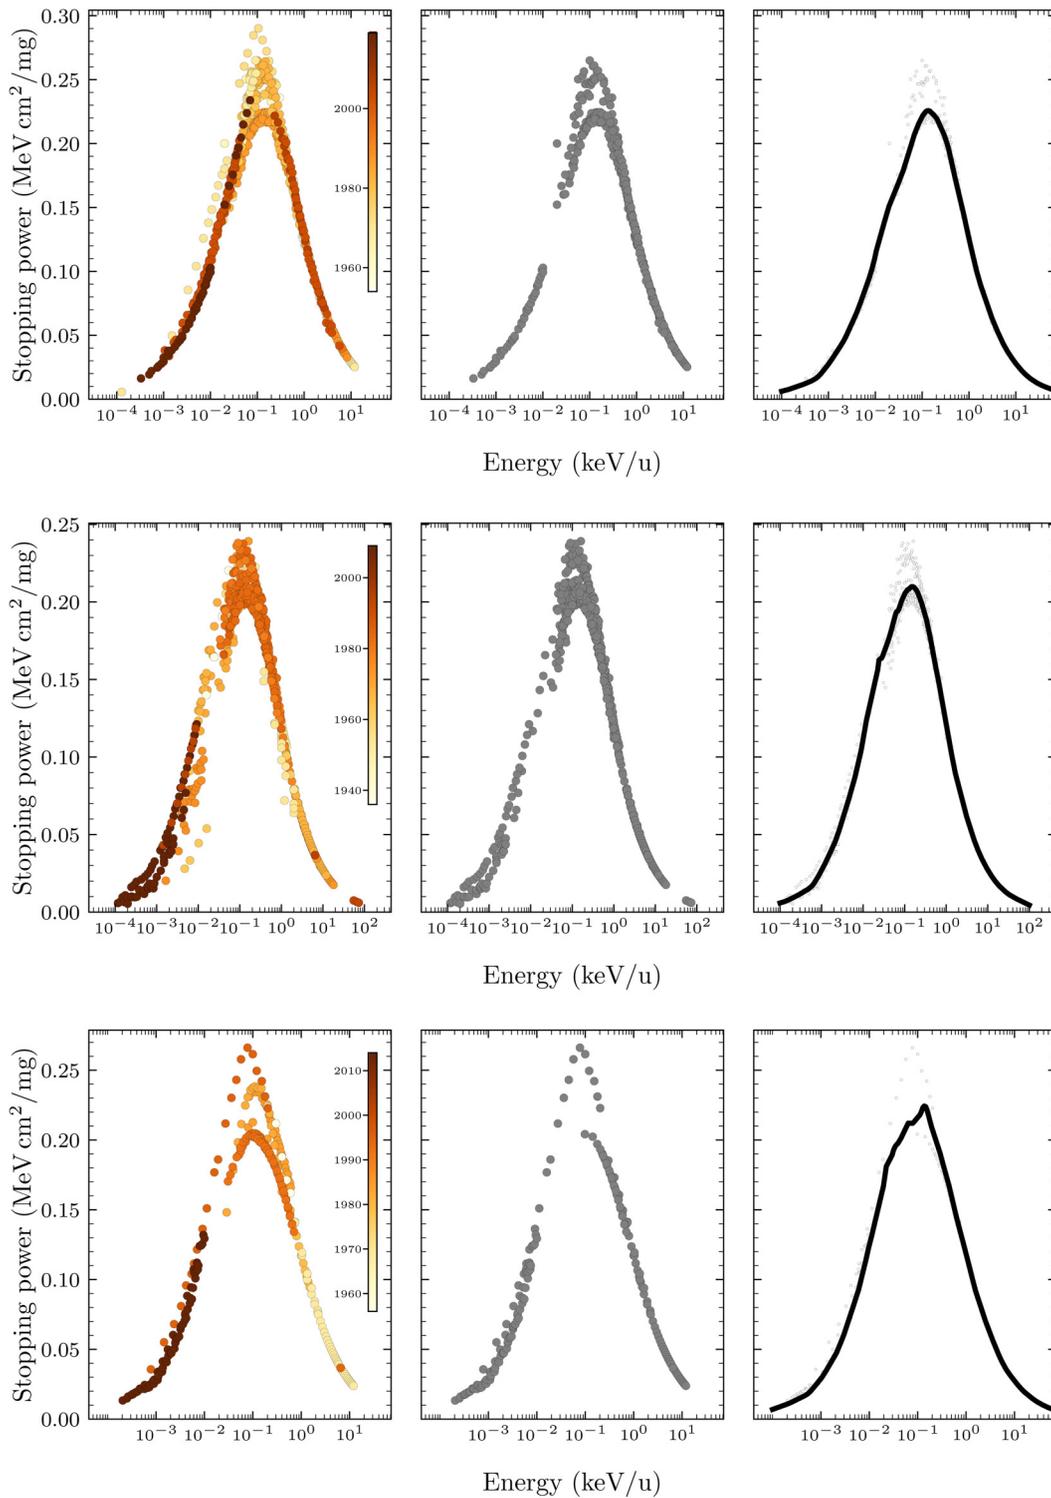
[Figure 5](#) shows stopping power cross sections of He colliding with Ne (top), Si (middle), and Cu (bottom). In some cases, the filtering procedure (from the left to the center figures) is hardly noticeable; in others, the year-of-publication criterion in the cleaning algorithm plays an important role. Examples of these cases are He in Ne and He in Cu, respectively.

For projectiles heavier than He, the amount of experimental data in the database (and generally measured) is scarcer than for H and He. For example, the cases Si in Si, O in C, and C in Al are displayed in [Fig. 6](#). Although these collisional systems have fewer data points, the filtering procedure was performed without further consideration. The outcome in most cases is strongly noticeable, as the center subplots from [Fig. 6](#) illustrate.

It was not apparent beforehand how severe the filtering had to be to proceed with the following supervised method. Hard filtering (resulting in clean and smooth data) seems to lead the neural network to overfit the input values while allowing some “noise” to prevent this problem. The present clustering-based algorithm reduced the original 36 000 data points to 28 000 values. These filtered measurements constitute the input database and are used to train a deep neural network, which will be described in [Sec. III](#).



**FIG. 3.** Left: Experimental results for stopping power cross sections of collisional systems with Au as targets. The colors indicate the year of publication of the data. Center: Filtered data, resulting from the cleaning procedure explained in the text. Right: Predicted data from the neural network. Top: H in Au. Middle: He in Au. Bottom: O in Au.



**FIG. 4.** Left: Experimental results for stopping power cross sections for H projectiles. The colors indicate the year of publication of the data. Center: Filtered data. Right: Predicted data from the neural network. Top: H in Ni. Middle: H in Cu. Bottom: H in Zn.

### III. STOPPING POWER PREDICTIONS

Calculating the stopping power of an ion when interacting with a target system involves numerous processes and parameters, which makes this a challenging problem. The different theoretical models developed to describe the experimental data accurately have not yet achieved total success. Semiempirical methods are generally needed to fit the known results and provide recommended values for multipurpose simulations and applications. Moreover, the experimental data also suffer from many discrepancies, requiring the introduction of auxiliary methods to classify and select the results. Considering the copious amount of experimental values in the IAEA database, it constitutes an ideal scenario where machine learning methods can help handle the complexities of this multi-process phenomenon and even discover hidden relations in the data.

At the beginning of this project, two different papers following this approach were published. Parfitt and Jackman<sup>21</sup> trained a random forest regression algorithm using thousands of measurements collected in Paul's stopping power database (with data up to 2015). Based on extensive evaluations through  $k$ -fold cross-validation against several error metrics, they demonstrated that their model could fit the training data and make low-error predictions on unseen test data. Using the same database, Guo *et al.*<sup>22</sup> trained a deep neural network to reproduce the experimental results for elementary targets with good accuracy.

#### A. Model

While designing our neural network, a thorough analysis was conducted to establish the best hyperparameters to use. This study includes not only the architectural design of the NN but also the definition of the loss function to minimize, the parameters involved in the minimization process, the learning rate, weight decay, and others. We found the best neural-network design consisting of a grid of five fully connected hidden layers, with  $10 \times 24 \times 32 \times 24 \times 10$  elements, a leaky-ReLu activation function for every layer, and the adaptive moment estimation (Adam) optimizer as the minimization algorithm. The dropout rates are 0.2 for the input and output layers and 0.5 for every hidden layer. The other parameters for the training model are the learning rate  $\alpha = 0.001$ , the batch size  $b = 64$ , and the weight decay  $\lambda = 10^{-10}$ . We employed 300 epochs with an early stopping step of 50. The re-parametrization trick<sup>23</sup> was used to increase the convergence speed.

It is worth mentioning the following technical detail. The stopping power must monotonically decrease toward low energies. However, for some collisional systems, an unphysical increase of the predicted values can appear toward small energy values. We tried adding minimal fictitious stopping power values at very low energies, but as expected, the training became extremely unstable due to its MAPE component. The correct behavior is accomplished in the final model by a simple trick: removing the bias parameters from the first linear layer.

An additional study focused on allocating the data for the learning process. From the cleaned data (28 000 values), we separated 5% (1400 results) for the test set. These data did not participate in the training process. The remaining 95% of the data, called the learning set, was partitioned into five folds. Four folds

comprised the training set, while the fifth fold constituted the validation set. Using this design, we trained the neural network until the loss function successfully converged. Then, the weights of every layer are kept to estimate the stopping power cross sections in a subsequent stage. This procedure is carried out iteratively by rotating the training-validation folds. As a result, we obtain five sets of weights and, therefore, five different estimations for any input entry. The final predictions are attained by averaging the results obtained from the five NN weights, also providing a rough estimation of the uncertainty present in the calculation.

As explained in Parfitt and Jackman,<sup>21</sup> different error metrics can be used to quantify the performance of the model in terms of the predicted values  $y_{\text{pred}}$  against the true experimental values  $y_{\text{true}}$ . For comparisons with other methods, we use the mean absolute percentage error loss function, or MAPE, given by

$$\text{MAPE} \equiv \frac{100}{n} \sum \left| \frac{y_{\text{true}} - y_{\text{pred}}}{y_{\text{true}}} \right|. \quad (4)$$

However, for the training process, the loss function is better defined by adding a small contribution of another metric to the MAPE. This prevents the training from featuring instabilities, which can appear at small stopping power values. We defined the loss function as a linear combination of two distinct metrics,

$$L \equiv \text{MAPE} + \frac{\text{MSE}}{\beta}, \quad (5)$$

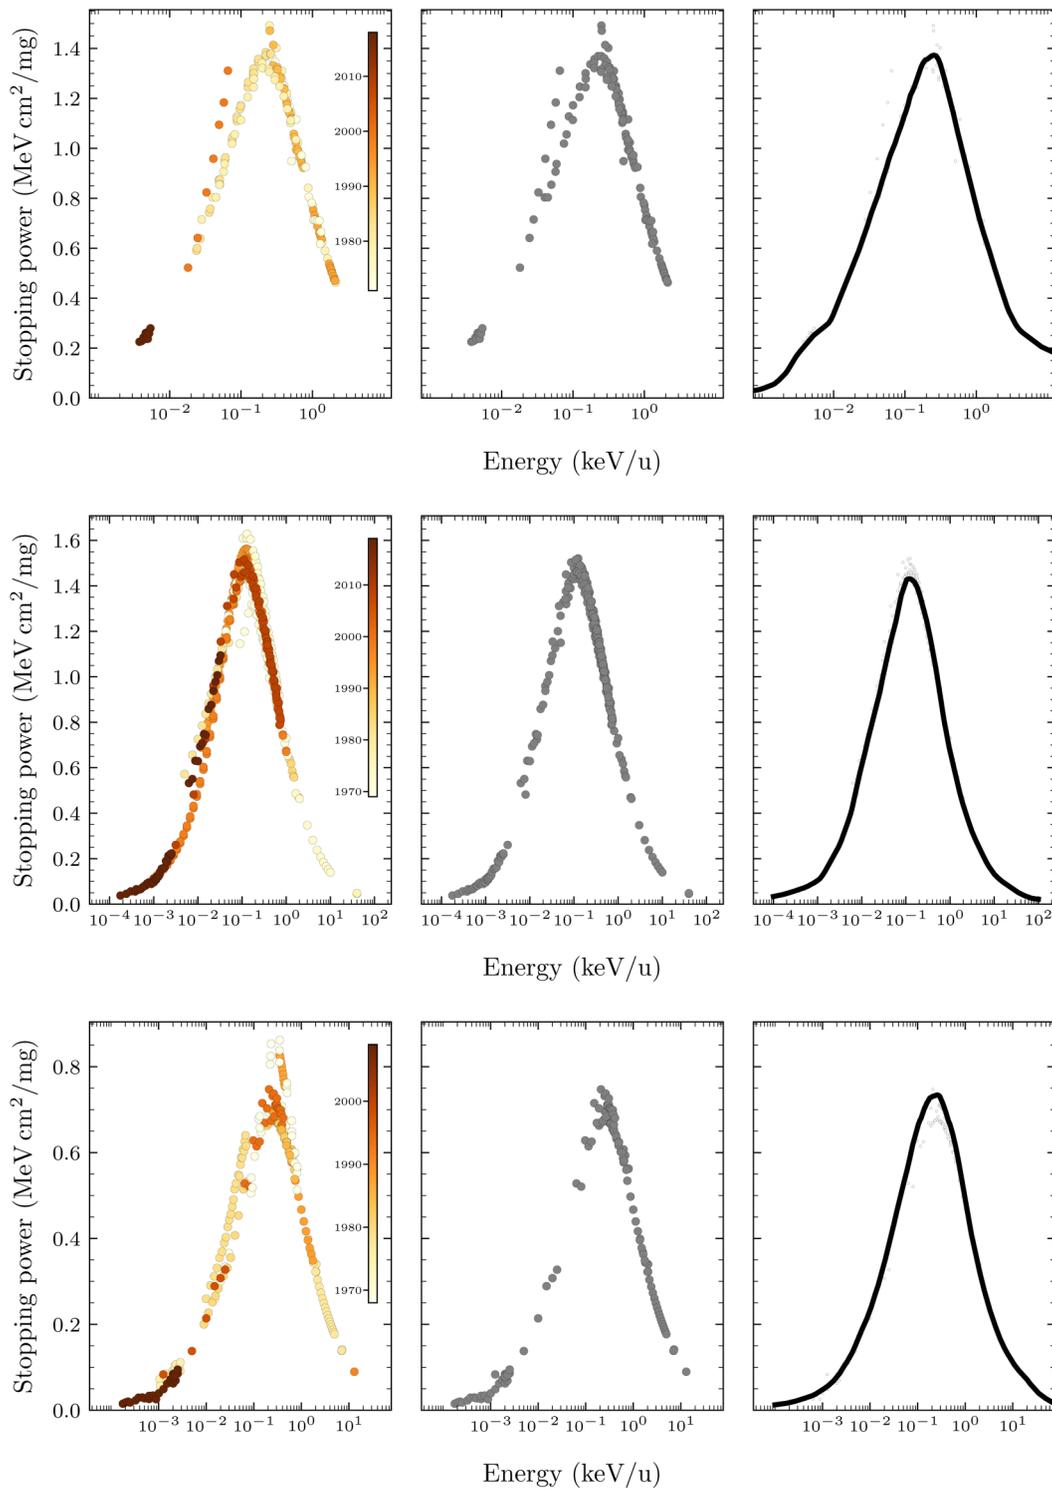
where MSE is the mean-squared error defined by

$$\text{MSE} = \frac{\sum (y_{\text{true}} - y_{\text{pred}})^2}{n}, \quad (6)$$

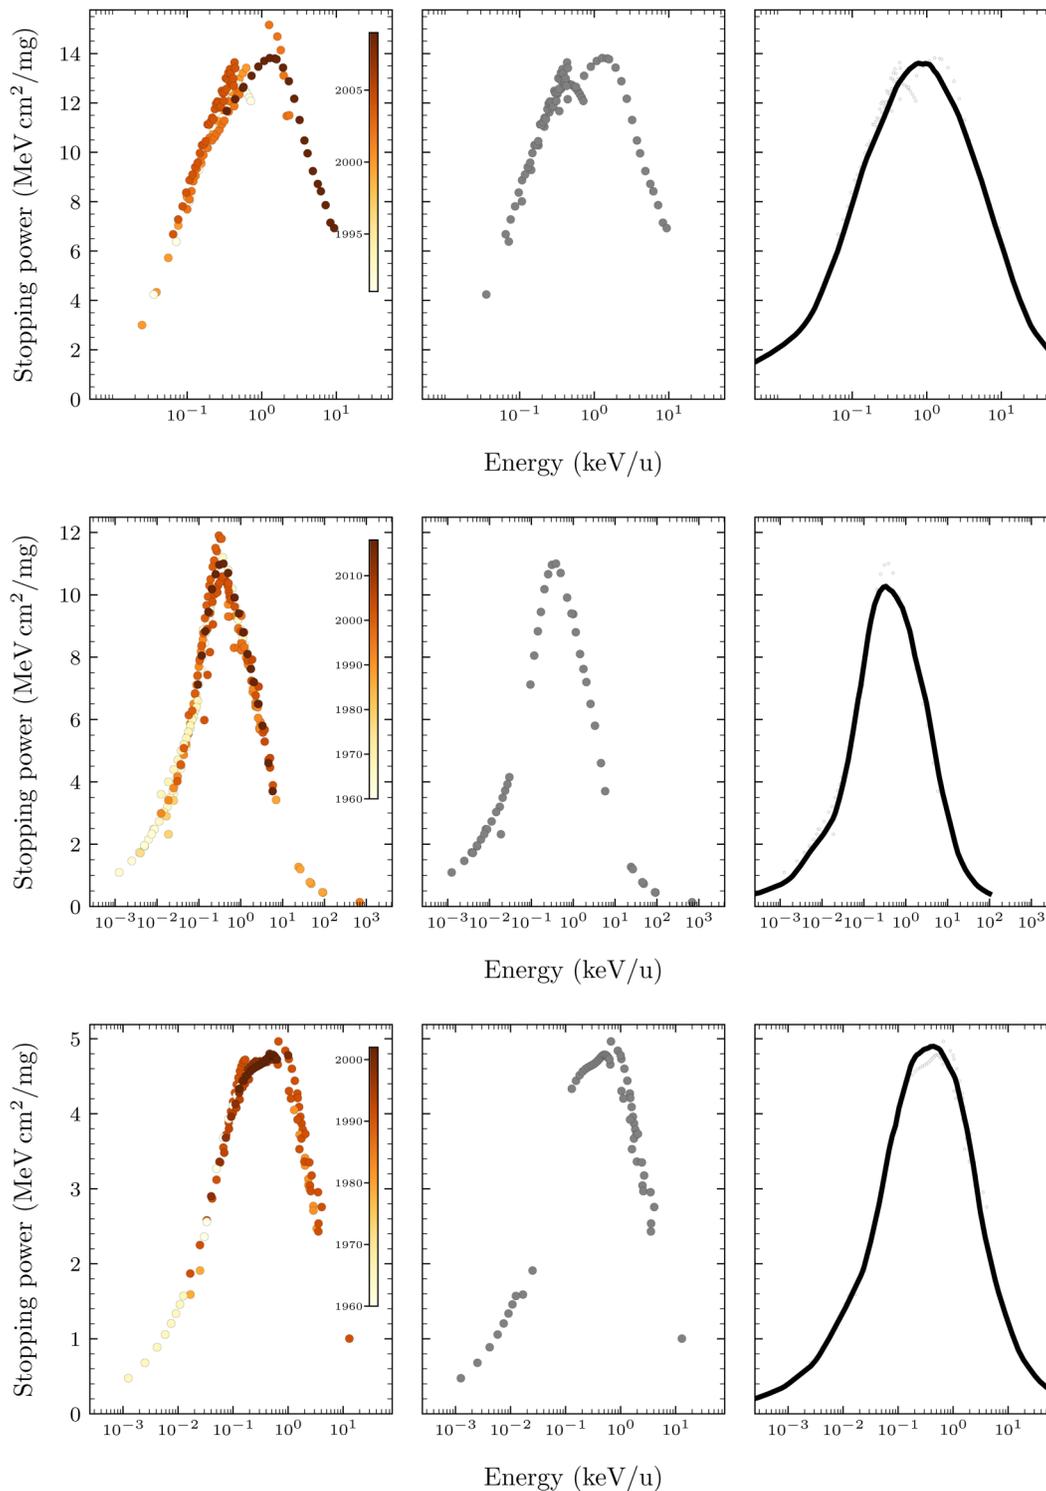
and  $\beta$  acts as a scaling factor between different units and is set to 100. Figure 7 shows the learning curve of four training sets, and we can observe that the learning procedure reaches a global minimum for the loss function. This behavior demonstrates that our NN model does not suffer from bias problems. Also, the fact that the validation set follows the same learning curve indicates that the data are not overfitted, which means that the neural-network parameters have been correctly selected, avoiding variance problems.

#### B. Features

Another aspect that must be taken into account when designing the NN model concerns the features to be considered from the input data. The network was trained with various sets of input features, and we concluded that it is possible to accurately determine the electronic stopping power of atoms with only five input features: the atomic mass of the incident ion, the atomic mass of the target, the atomic numbers of both the projectile and the target, and the incident projectile energy (per unit mass, i.e., keV/amu). The corresponding MAPE values are summarized in Table I. The first improvement is introduced by considering the logarithm of the incident energy in place of the actual value. We also added the first ionization energy of the target, obtaining a further



**FIG. 5.** Left: Experimental results for stopping power cross sections for He projectiles. The colors indicate the year of publication of the data. Center: Filtered data. Right: Predicted data from the neural network. Top: He in Ne. Middle: He in Si. Bottom: He in Cu.



**FIG. 6.** Left: Experimental results for stopping power cross sections. The colors indicate the year of publication of the data. Middle: Filtered data. Right: Predicted data from the neural network. Top: Si in Si. Middle: O in C. Bottom: C in Al.

improvement in the MAPE values. Encouraged by this, we tested adding other features, such as the second ionization limit of the target, the ionization energy of the ion, and the electronegativities of both of them, resulting in poorer outcomes with higher error values.

#### IV. RESULTS

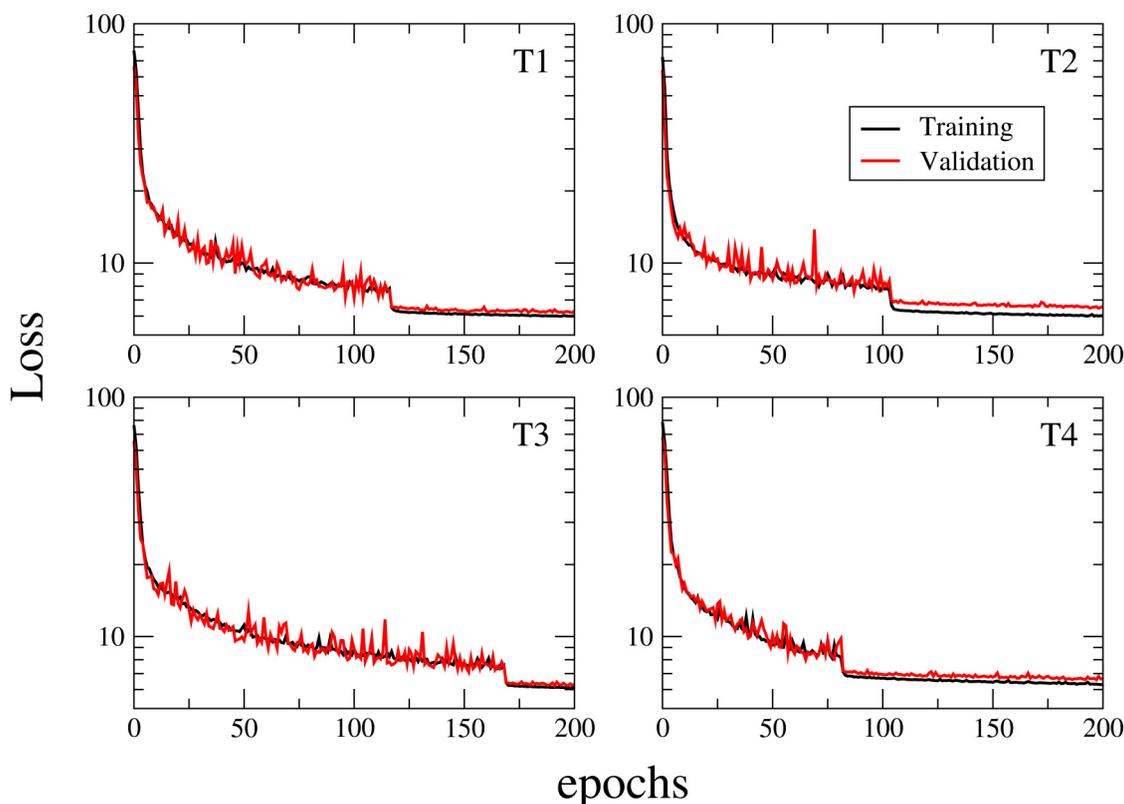
This section presents and examines the results obtained with the ESPNN model, which has been described in Sec. III A. The stopping power cross section predictions result from the input forward propagation along the neural network. In the subplots at the right of Figs. 3–6, we show the ESPNN predictions for the collisional systems examined in Sec. II C. For these ion–target combinations, the model replicates very well the experimental results. It is noteworthy that although the predictions have been obtained individually for each projectile energy value, the model’s results are given by smooth curves. The stopping power cross sections shown in Fig. 3 for multiple collisional systems with Au as the target tend to follow the trend given by the latest observations. On the contrary, the ESPNN results illustrated in Fig. 4 for the collisional systems with H as projectiles replicate the shape of the more statistically representing data. The smoothness of the ESPNN stopping power cross sections can be understood as a consequence of using

**TABLE I.** Error values evaluated on the cross-validation dataset, resulting with different input features considered in the model training. Boldface denotes the selected set of input features used in the model.

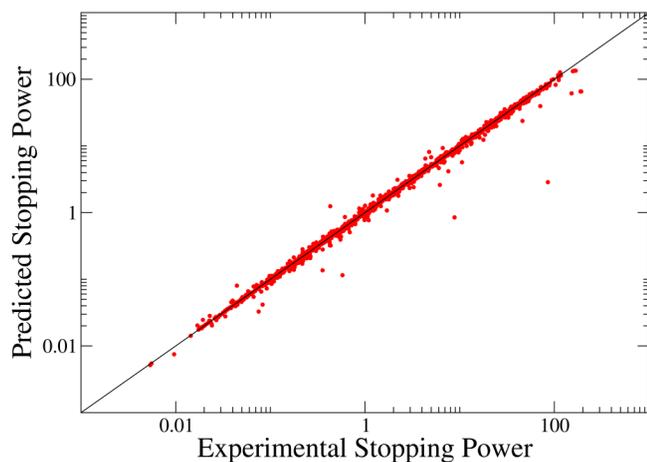
Features	MAPE (%)
Default: $Z_p, m_p, Z_t, m_t, E$	5.76
$E \rightarrow \log E$	5.47
+ first ionization (target)	<b>5.07</b>
+ first + second ionization (target)	14.9
+ first ionization (target) + first ionization (projectile)	16.1
+ first ionization and electronegativity (target)	5.11
+ first ionization (target) + electronegativity (projectile)	23.8

a reduced number of features in the training process, which also avoids overfitting and any renormalization of the input variables. Nonetheless, having a result of a soft curve is somewhat remarkable since the experimental input data may present significant spreading in some energy regions.

To verify the transferability of the model, we calculated the stopping power cross sections for the experimental data points initially excluded from the training procedure: the test set. Even though these values were not included in the training (or



**FIG. 7.** Learning curve for the training and validation datasets for four different folds.



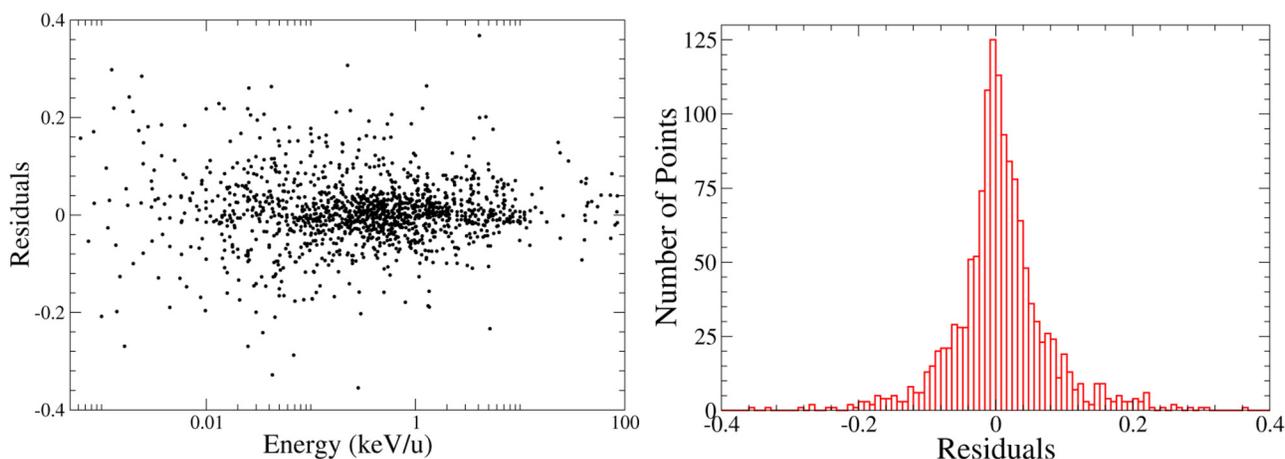
**FIG. 8.** Predicted ESPNN values for the stopping power cross sections, compared with the experimental results. All the results correspond to the test dataset, never seen before by the model.

validation) sets, our model can reproduce these results with high accuracy, as shown in Fig. 8.

We can further examine the performance of the ESPNN on the test set by analyzing the residuals given by

$$R \equiv \frac{y_{\text{true}} - y_{\text{pred}}}{y_{\text{true}}}. \quad (7)$$

In Fig. 9, we show the test set residuals over the experimental energy range (left panel) and the frequency distribution of these values (right panel). We can estimate a mean error value of 5% from this figure. Because the experimental data are scattered over a



**FIG. 9.** Error of the predicted ESPNN values for the test dataset. Left: residual values. Right: Histogram of the prediction errors.

**TABLE II.** Comparisons between the MAPE values reported by SRIM,<sup>14</sup> the random forest (RF) algorithm from Parfitt and Jackman,<sup>21</sup> the deep learning (DL) algorithm from Guo *et al.*,<sup>22</sup> and the present calculations (ESPNN). Left columns: error values in the training set. Right columns: error values in the test set.

Ions	Train			Test	
	SRIM	RF	DL	RF	ESPNN
H	4.0	2.4	5.6	12	<b>4.5</b>
He	3.9	2.1	6.0	9.1	<b>4.1</b>
Li	4.8	2.1	4.2	25	<b>7.7</b>
Be-U	5.8	2.0	7.0	12	<b>7.1</b>
All	4.6	2.1	5.7	23	<b>5.7</b>

broad range of uncertainties and discrepancies, the present model's accuracy can be considered exceptional.

For a quantitative comparison of our model with other approaches, we show in Table II the MAPE values reported by SRIM,<sup>14</sup> the random forest (RF) algorithm from Parfitt and Jackman,<sup>21</sup> the deep learning (DL) algorithm from Guo *et al.*,<sup>22</sup> and the present ESPNN calculations. The table does not provide a rigorous comparison between these methods because it displays the values as reported by the respective authors, and different numbers of data points were considered in the final MAPE evaluation. For that reason, we separated the table into two parts, showing on the left the published errors for the training set and on the right the MAPE in the test set. However, we can only make a proper comparison with the RF calculations reported by Parfitt and Jackman. Although the size and characteristics of this work's input and test sets were chosen differently than ours, these authors addressed the training and an unseen test set separately. The MAPE values in Table II provide a clear picture of the various code's performance. Furthermore, the ESPNN results on the unseen test set are excellent. A few details are worth mentioning; SRIM is a semiempirical code that considers all the experimental data published until 2013;

**TABLE III.** Predictive MAPE of the semi-empirical SRIM code<sup>14</sup> and the neural-network ESPNN model trained only with data collected before 2013. The first columns show the error predictions for all the experimental values reported after 2013. The second includes all the values reported after 2015. The last columns show the error of the predictions for the experimental results produced after 2019.

Ions	>2013		>2015		>2017		>2019	
	ESPNN	SRIM	ESPNN	SRIM	ESPNN	SRIM	ESPNN	SRIM
H	7.0	19.2	4.6	15.3	6.6	13.3	3.5	7.1
He	8.4	10.6	9.3	10.1	9.3	10.1	5.5	8.3
Be-U	6.2	6.6	5.4	6.5	5.3	6.3		
All	7.0	11.4	6.0	9.8	6.8	8.9	4.1	7.4

therefore, the respective errors do not necessarily correspond to its predictive capability on later measurements. Inversely, RF describes the training data very well but overfits; the high MAPE values for the test set may be an indication of this.

Because of the importance and the widespread usage of SRIM, we devised a detailed experiment to compare our model's calculations. Attempting to make a fair comparison between our predictions and SRIM, we conducted the following procedure: we trained the neural network only with results published before 2013. This means that we applied our DBSCAN-based heuristic to all the data published up to 2013. The remaining experimental values (published after 2013) were set aside and left unclear to avoid data leakage. The excluded set (the new test set) was further separated into different groups in an iterative fashion. First, we considered all the experimental results published after 2013, and next, all the publications that appeared after 2015 and successively reduced the test set to include only the most recent publications. Noteworthy, this neural network was trained only once, with the data collected in the IAEA database until 2013. In this way, we can compare the prediction power of both methods on an equal footing. The results are summarized in Table III. This table shows that the ESPNN prediction power is consistently better than SRIM's in all the cases. Also, comparing these results with the MAPE values obtained with the complete data training, we observe that the prediction power is improved, as expected, as new values are included in the data set.

## V. CONCLUSIONS

In this work, the difficult task of predicting the electronic stopping power cross sections for atomic targets is addressed using two ML methods. First, a clustering-based algorithm was developed to eliminate automatically suspicious or erroneous experimental values. The method implements the DBSCAN algorithm to detect outliers and clusters with similar data. A multi-step cleaning algorithm analyses these results and implements three specific criteria. The algorithm was prepared to filter about 20% of the 36 000 experimental atomic results in the IAEA stopping power database.

The filtered data were used to train a deep neural network, constituting the second ML method. The neural network was trained to accurately reproduce the training data, producing MAPE results of less than 6%. Considering the widespread of the experimental values, it represents excellent performance. The same MAPE was obtained for the test set, which constituted data not

included in the training procedure. These results demonstrate the excellent prediction power of our model.

Moreover, an open-access ESPNN code with the forward propagation model has been published in a public repository and is now available. This algorithm allows computing the stopping power cross section of any collisional system effortlessly and quickly. Certainly, the model will be improved in the future with the addition of newer experimental data and stopping cross sections of complex molecules.

## ACKNOWLEDGMENTS

The following institutions of Argentina financially support this research: the CONICET (No. PIP11220200102421CO), the ANPCyT (No. PICT-2020-SERIE A-01931), and the University of Buenos Aires (Project No. 20020170100727BA).

## AUTHOR DECLARATIONS

### Conflict of Interest

The authors have no conflicts to disclose.

### Author Contributions

**F. Bivort Haiek:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Methodology (equal); Software (lead); Writing – review & editing (equal). **A. M. P. Mendez:** Software (equal); Visualization (equal); Writing – review & editing (equal). **C. C. Montanari:** Conceptualization (equal); Data curation (lead); Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Validation (equal); Writing – review & editing (equal). **D. M. Mitnik:** Conceptualization (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Software (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (lead).

## DATA AVAILABILITY

The data that support the findings of this study are openly available in GitHub at <https://github.com/ale-mendez/ESPNN>, Ref. 24, and in PyPI at <https://pypi.org/project/ESPNN/>, Ref. 25.

APPENDIX A: MACHINE LEARNING FILTERING  
PROCEDURE

---



---

```

1: procedure CLUSTERS AND OUTLIERS
2: Input:
   1.  $\epsilon$  (radius of the neighborhood)
   2.  $N_{min}$  (number of reachable points)
3: Run DBSCAN to select different clusters and outliers.
4: Output:
   1.  $N_{clust}$  (number of clusters)
   2.  $N_{outl}$  (number of noise points)
5: end procedure CLUSTERS AND OUTLIERS
6:
7: procedure Drop Publications
8: for list of publications  $P_i$  do
9:
10: procedure Check Newer Publications
11:   Count  $n_i$  (number of experimental results  $P_i$ )
12:   Define  $\Delta E_i$  (energy range covered in  $P_i$ ).
13:   for list of publications  $P_j^n$  newerthan  $P_i$  do
14:     Define  $\Delta E_{ji}^n \equiv \Delta E_j^n \cap \Delta E_i$ .
15:   end for
16:   Define  $\Delta E_{ji}^{new} \equiv \bigcup_j \Delta E_{ji}^n$ 
17:   if  $\frac{\Delta E_{ji}^{new}}{\Delta E_i} \leq \sigma_{\Delta}$ : KEEP publication  $P_i$  and BREAK
18: end procedure CHECK NEWER PUBLICATIONS
19:
20: procedure ISOLATED RESULTS
21:   Count  $N_{outl}^i$  (number of outliers in  $P_i$ )
22:   if  $\frac{N_{outl}^i}{n_i} > \sigma_{outl}$ : DROP publication  $P_i$ 
23: end procedure ISOLATED RESULTS
24:
25: procedure TEST CLUSTER
26:   Identify  $C^i$  (biggest cluster in  $P_i$ )
27:   Count  $l^i$  (number of  $n_i$  values from  $P_i \in C^i$ )
28:   Count  $t^i$  (total number of results  $\sum n_j \in C^i$ )
29:   if  $\frac{l^i}{t^i} > \sigma_{clu}$ : DROP publication  $P_i$ 
30: end procedure TEST CLUSTERS
31:
32: end for
33: end procedure DROP PUBLICATIONS

```

---



---

## APPENDIX B: INSTRUCTIONS FOR RUNNING ESPNN

The code is public and can be used remotely or locally. Detailed instructions can be found at

<https://pypi.org/project/ESPNN/>.

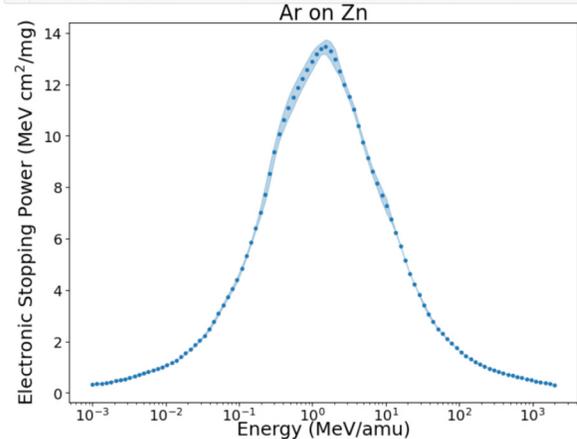
For remote use, the ESPNN code can run on the Google Colab platform (see the url address in the PyPI page). For a local installation, the code can be installed via pip:

```
pip install ESPNN
```

or by downloading or cloning the repository at

<https://github.com/ale-mendez/ESPNN>.

```
In[14]: 1 ESPNN.run_NN(projectile='Ar', target='Zn', emin=0.001, emax=2000, npoints=100)
```



**FIG. 10.** Example of an output figure: Ar in Zn. The stopping power cross sections are predicted at  $n_{\text{points}} = 100$  points, and the estimated variance is plotted in light blue.

Once the code is installed, the execution is extremely simple. For example, if the He (projectile) and Au (target) collision is required, the user can type from the terminal

```
python -m ESPNN He Au.
```

To run the code using a Jupyter Notebook, only these two lines are required:

```
import ESPNN
ESPNN.run_NN(projectile="He", target="Au")
```

There are other optional arguments that can be modified, such as the minimum and maximum energy values, the number of energy points in the grid, the path to the output file, etc. For information about the options,

```
python -m ESPNN -h
```

As an example, we show the results for Ar in Zn. An output file called ArZn\_prediction.dat is generated. This file has three columns: Energy (MeV/amu), Stopping power (MeV cm<sup>2</sup>/mg), and variance (MeV cm<sup>2</sup>/mg). Figure 10 is generated on the notebook.

## REFERENCES

- <sup>1</sup>See <https://www-nds.iaea.org/stopping/> for “Electronic Stopping Power of Matter for Ions: Graphs, Data, Comments and Programs.”
- <sup>2</sup>H. Paul, D. Semrad, and A. Seilinger, *Nucl. Instrum. Methods Phys. Res. B* **61**, 261 (1991).
- <sup>3</sup>H. Paul and A. Schinner, *At. Data. Nucl. Data Tables* **85**, 377 (2003).
- <sup>4</sup>H. Paul, *AIP Conf. Proc.* **1525**, 295 (2013).
- <sup>5</sup>H. Paul, *AIP Conf. Proc.* **1525**, 309 (2013).
- <sup>6</sup>H. Paul, “Stopping Power for Light Ions”; see <https://www-nds.iaea.org/stopping/>; original webpage by H. Paul; see <http://www.exphys.jku.at/stopping/>.
- <sup>7</sup>For a list of Paul’s publications on stopping power, see <https://www-nds.iaea.org/stopping/stoppingdocu.html>.
- <sup>8</sup>C. C. Montanari and P. Dimitriou, *Nucl. Instrum. Methods Phys. Res. B* **408**, 50 (2017).

- <sup>9</sup>P. Sigmund, *Particle Penetration and Radiation Effects. Vol. 1: General Aspects and Stopping of Swift Point Charges* (Springer-Verlag, Berlin, 2006).
- <sup>10</sup>P. Sigmund, *Particle Penetration and Radiation Effects. Vol. 2: Penetration of Atomic and Molecular Ions* (Springer International Publishing, Switzerland, 2014).
- <sup>11</sup>P. Sigmund and A. Schinner, *Nucl. Instrum. Methods Phys. Res. B* **382**, 15 (2016).
- <sup>12</sup>"Stopping Power of Matter for Ions. Computer Programs"; see <https://www-nds.iaea.org/stopping/stoppingprog.html>.
- <sup>13</sup>J. F. Ziegler, J. P. Biersack, and U. Littmark, *The Stopping and Range of Ions in Solids* (Pergamon Press, 1985); see <http://www.srim.org/>.
- <sup>14</sup>J. F. Ziegler, M. D. Ziegler, and J. P. Biersack, *Nucl. Instrum. Methods Phys. Res. B* **268**, 1818 (2010).
- <sup>15</sup>N. P. Barradas, A. Bergmaier, K. Mizohata, M. Msimanga, J. Räisänen, T. Sajavaara, and A. Simon, *Nucl. Instrum. Methods Phys. Res. B* **360**, 90 (2015).
- <sup>16</sup>T. Materna, E. Berthoumieux, Q. Deshayes, D. Doré, M. Kebbiri, A. Letourneau, L. Thulliez, Y. H. Kim, U. Köster, and X. Ledoux, *Nucl. Instrum. Methods Phys. Res. B* **505**, 1 (2021).
- <sup>17</sup>M. V. Moro, P. Bauer, and D. Primetzhofer, *Phys. Rev. A* **102**, 022808 (2020).
- <sup>18</sup>F. F. Selau, H. Trombini, G. G. Marmitt, A. M. H. De Andrade, J. Morais, P. L. Grande, I. Alencar, M. Vos, and R. Heller, *Phys. Rev. A* **102**, 032812 (2020).
- <sup>19</sup>M. Ester, H. P. Kriegel, J. Sander, and X. Xu, *Knowl. Discov. Datab.* **96**, 226 (1996).
- <sup>20</sup>E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, *ACM Trans. Database Syst.* **42**(3), 19 (2017).
- <sup>21</sup>W. A. Parfitt and R. B. Jackman, *Nucl. Instrum. Methods Phys. Res. B* **478**, 21 (2020).
- <sup>22</sup>X. Guo, H. Wang, C. L. Shijun Zhao, K. Jin, and J. Xue, *Chin. Phys. B* **31**, 073402 (2022).
- <sup>23</sup>T. Salimans and D. P. Kingma, in *Advances in Neural Information Processing Systems*, edited by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), Vol. 29.
- <sup>24</sup>See <https://github.com/ale-mendez/ESPNN> for ESPNN (Electronic Stopping Power Neural Network).
- <sup>25</sup>See <https://pypi.org/project/ESPNN> for ESPNN (Electronic Stopping Power Neural Network).