Deep learning analysis of deeply virtual exclusive photoproduction

Jake Grigsby,^{*} Brandon Kriesten,[†] Joshua Hoskins[‡], and Simonetta Liuti[§]
Department of Physics, University of Virginia, Charlottesville, Virginia 22904, USA

Peter Alonzi^{||}
School of Data Science, University of Virginia, Charlottesville, Virginia 22904, USA

Matthias Burkardt[¶]
*New Mexico State University, Department of Physics,
 Box 30001 MSC 3D, Las Cruces, New Mexico 88003, USA*

(Received 15 December 2020; accepted 22 March 2021; published 1 July 2021)

We present a machine learning based approach to the cross section and asymmetries for deeply virtual Compton scattering from an unpolarized proton target using both an unpolarized and polarized electron beam. Machine learning methods are needed to study and eventually interpret the outcome of deeply virtual exclusive experiments since these reactions are characterized by a complex final state with a larger number of kinematic variables and observables, exponentially increasing the difficulty of quantitative analyses. Our deep neural network (FemtoNet) uncovers emergent features in the data and learns an accurate approximation of the cross section that outperforms standard baselines. FemtoNet reveals that the predictions in the unpolarized case systematically show a smaller relative median error than the polarized that can be ascribed to the presence of the Bethe-Heitler process. It also suggests that the t dependence can be more easily extrapolated than for the other variables, namely the skewness, ξ and four-momentum transfer, Q^2 . Our approach is fully scalable and will be capable of handling larger datasets as they are released from future experiments.

DOI: [10.1103/PhysRevD.104.016001](https://doi.org/10.1103/PhysRevD.104.016001)

I. INTRODUCTION

Artificial intelligence (AI) methods including machine learning (ML), deep learning (DL) and neural networks (NN) have been increasingly applied to attack several of the outstanding questions on strongly interacting systems in quantum chromodynamics (QCD), ranging from both the sign problem (see [1–3] and references therein), and the inverse problem in lattice QCD, to extracting the behavior of parton distributions functions (PDFs) from data. However, the systematic application of AI to problems in hadronic physics can still be considered to be in its initial stages, which involves exploring a variety of

methodologies [4]. One important aspect that is still being developed concerns the treatment of uncertainty, in particular, studying the accuracy of AI based predictions within a well defined statistical approach and with precise benchmarking.

In this paper we address the recently developed science of nuclear femtography, which aims at reconstructing the spatial structure of hadrons, or nucleon tomography, using information from deeply virtual electron proton exclusive photoproduction, $ep \rightarrow e'p'\gamma$, and various other related reactions. The open questions that we address concern, in particular, the choice of a ML method that is suitable to attack the issues specific to nuclear femtography, including the impact of noise and experimental uncertainty.

Deeply virtual exclusive processes have been identified as the cleanest experimental probe of the 3D structure of the proton as well as its mechanical properties including the quark and gluon distributions of angular momentum, pressure, and shear forces (see reviews in [5–8]). Information on these distributions is encoded in the generalized parton distributions (GPDs) which are functions of the longitudinal momentum fraction, x , and of the longitudinal momentum transfer, ξ (skewness parameter) and four momentum transfer squared, $t = \Delta^2$, at a given

^{*}jcg6dn@virginia.edu
[†]btk8bh@virginia.edu
[‡]jhoskins@jlab.org
[§]sl4y@virginia.edu
^{||}lpa2a@virginia.edu
[¶]burkardt@nmsu.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

scale determined by the photon virtuality, Q^2 . GPDs can be defined for all the different polarization configurations of the struck quark/gluon and proton. Their Fourier transform with respect to the transverse component of Δ produces faithful images of the proton constituents structure on the transverse plane.

In exclusive photoproduction we access GPDs from the matrix elements for deeply virtual Compton scattering (DVCS), where the photon is produced at the proton vertex. A competing background process given by the Bethe-Heitler (BH) reaction is also present, where the photon is emitted from the electron and the matrix elements measure the proton elastic form factors. Note that for this prototypical process, the cross section is a function of four independent kinematic variables: the angle between the lepton and hadron planes, ϕ , the scale, Q^2 , the skewness, ξ , which is related to Bjorken x , and t ; GPDs for all quark flavors appear in the cross section embedded in convolutions over x known as Compton form factors (CFFs), with complex kernels depending on (x, ξ, Q^2) . Several CFF combinations corresponding to various quark-proton polarization configurations appear simultaneously in the cross section terms for each beam-target polarization configuration [9]. The complexity associated with the large number of kinematic variables, and the complex structure of the cross section makes the problem of extracting tomographic images of quark and gluon structure of the proton virtually impossible to solve with traditional methods; for high precision femtography, which is required to obtain proton images, we need to develop more sophisticated analyses. The success of ML methodologies in modeling complex natural phenomenon make this a prime choice for GPD extraction.

The ultimate goal is development of a tool set for the extraction of GPD from experimental data; however, as a first step, we develop a framework for the extraction of the scattering cross section. The results of our analysis show that the network can effectively generalize in t , even in the regions with no data, but is less effective in ξ , or Q^2 . These results are not only beneficial to the physics community but provide interesting overlays with the data science community. One important example is discussion of handling experimental error bars found in Sec. III B 1, which is ubiquitous to physics analyses but less commonly considered in building machine learning models.

We use industry standard ML techniques to fit a cross section model based on currently available DVCS experimental data, allowing us to make efficient and accurate NN predictions interpolating between experimental data points across a wide kinematic range. By estimating model uncertainty, we can also make informed decisions about NN predictions well outside of the region defined by our data, extrapolating to kinematic regions unexplored by current experimental programs. Several of the practical challenges of fitting a sparse neural network, with

significant experimental uncertainty, are addressed in this paper. All of our software tools are open source and available for download [10].

This paper is organized as follows: in Sec. II we present the theoretical framework for DVCS and give background on deep learning and neural networks. In Sec. III we detail the specifics of our cross section network and training process. Then, in Sec. IV we use our networks to make predictions across a wide range of kinematic bins, comparing them against both experimental measurements and previous theory-driven models. Finally, our conclusions and outlook are presented in Sec. V.

II. BACKGROUND

A. DVCS theoretical framework

The cross section for $ep \rightarrow e'p'\gamma$ reads

$$\frac{d^5\sigma}{dx_{Bj}dQ^2dt|d\phi d\varphi_S} = \Gamma|T|^2. \quad (1)$$

where T is a superposition of the DVCS amplitude and its background process Bethe-Heitler (BH),

$$|T|^2 = |T_{\text{BH}} + T_{\text{DVCS}}|^2 = |T_{\text{BH}}|^2 + |T_{\text{DVCS}}|^2 + \mathcal{I} \quad (2)$$

$$\mathcal{I} = T_{\text{BH}}^* T_{\text{DVCS}} + T_{\text{DVCS}}^* T_{\text{BH}}. \quad (3)$$

The cross section is a differential form of five independent variables, the four momentum transfer squared between the initial (e) and final (e') electrons, Q^2 ; $x_{Bj} = Q^2/(2M\nu)$ where ν is the energy transfer between the initial and final electrons and M is the proton mass; t , the (negative) four momentum transfer squared between the initial and final proton; the angle ϕ between the lepton and hadron planes, and φ_S , the azimuthal angle of the transverse proton spin vector. We will consider in our analysis available fixed target data for an unpolarized target where the measured set of observables includes

$$E_e, Q^2, x_{Bj}, t, \phi,$$

E_e being the initial electron energy in the lab (the angle φ_S is integrated over since we are not considering scattering from a polarized target). The ultimate physics goal is to extract from experiment the QCD CFFs which appear in the DVCS matrix element. At leading order one can define eight CFFs, which correspond to the real and imaginary parts of the two vector and two axial vector current components, respectively. Before considering the full extraction of all eight CFFs, an important first step is to understand what is required to learn the form of the cross section. Our ML analysis yields model independent results for the cross section independent from the separation into the BH, DVCS and $\mathcal{I}$ components given in Eq. (2). Because

TABLE I. Experimental data used in our neural network analysis. No data from HERMES or COMPAS was used because it is integrated in φ and therefore inconsistent dimensionally with a majority of the JLAB DVCS data.

Experiment	Q^2 (GeV ²)	$t_{\min} - t_{\max}$ (GeV ²)	$x_{Bj}^{\min} - x_{Bj}^{\max}$	E_{beam} (GeV)	Polarization	Number of points	Reference
Hall B	1.11/3.77	-0.11/ - 0.45	0.13/0.48	5.75	UU/LU	3862	[11]
Hall A	1.45/2.38	-0.17/ - 0.37	0.34/0.40	5.75	UU/LU	936	[12]
Hall A	1.51/2.00	-0.18/ - 0.36	0.36	5.55	UU/LU	788	[13]
Hall A	2.71/8.51	-0.21/ - 1.28	0.34/0.61	4.49/10.99	UU/LU	2160	[14]

of the nature of the scattering experiments there are many more data on the ϕ dependence than on the other variables (see Table I in Sec. III A). In this paper we focus on the simplest case of unpolarized target scattering for which we also have the largest DVCS dataset.

B. Deep neural networks

A deep neural network (DNN) is an effective function approximation method consisting of several layers of computation, in which each layer computes a new representation of the previous layer's outputs. In the simplest case, each layer consists of an affine transformation of the previous layer's representation of the data, followed by an elementwise nonlinearity. For example, the output of the first layer is

$$\mathbf{x}_1 = z(\mathbf{W}_0 \mathbf{x} + \mathbf{b}_0), \quad (4)$$

where z is a nonlinear activation function applied to each element of its input vector. $\mathbf{W}$ is an $m \times n$ matrix where m is the dimension of the previous layer's representation and n is the dimension of the current layer's output. $\mathbf{b}$ is a vector of size n . The entire model can be written as a composition of functions

$$f_{\theta}(\mathbf{x}) = \mathbf{W}_k(z(\mathbf{W}_{k-1}(\dots(z(\mathbf{W}_0 \mathbf{x} + \mathbf{b}_0)))) + \mathbf{b}_{k-1})) + \mathbf{b}_k, \quad (5)$$

where k is the number of hidden layers in the model. Model “architectures” are descriptions of the number (k) and size (n) of a network's layers, as well as the choice of nonlinearity (z). θ is the set of all the entries in the transformation weights $\mathbf{W}$ and biases $\mathbf{b}$. These parameters are initialized randomly, and then adjusted by an optimization algorithm. The goal is to minimize some metric of prediction error—usually called the loss function. A standard choice for regression problems is the mean squared error (MSE):

$$\mathcal{L}_{\theta}^{\text{MSE}}(\mathbf{x}, \mathbf{y}) = \frac{1}{2N} \sum_{i=0}^{i=N} (\mathbf{y}_{(i)} - f_{\theta}(\mathbf{x})_{(i)})^2, \quad (6)$$

where y is the true value of our predicted variable, often called the “label” or “target” value. In our case, back-propagation is used to compute the partial derivatives of the

loss function on a batch of B samples with respect to each of the parameters in θ , and the resulting gradient vector is used to update the parameters via stochastic gradient descent (SGD):

$$\theta_{t+1} = \theta_t - \lambda \nabla_{\theta} \frac{1}{B} \sum_{i=0}^B \mathcal{L}_{\theta_i}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}), \quad (7)$$

where λ is a learning rate parameter. Modern optimizers [15,16] often dynamically adjust λ based on heuristics like momentum, where consecutive gradient steps in a similar direction increase the learning rate and rapid changes in direction push it towards zero.

III. TRAINING THE DEEP NETWORK

A. Experimental data

We compiled a dataset from several DVCS experiments [11–14], covering the range of kinematic bins shown in Fig. 1. We discarded datasets that listed no explicit φ dependence. We used a common format for collecting and recording our datasets such that they list the kinematic dependence of the cross section points, the ϕ dependence, the polarization of the observable, and the statistical and systematic errors.

Each entry consists of the input kinematic variables (x_{Bj} , t , Q^2 , E_b , φ) as well as the value and experimental error of the cross section. In total, we are working with 3,862 unpolarized and 3,884 polarized data points. Because these measurements are typically taken by sweeping across φ values in a fixed kinematic bin, most points are unique only in φ , meaning there are relatively few values of x_{Bj} , t , Q^2 , E_b . This presents an interesting challenge for ML approaches, which become much more capable of predicting across φ values than they are of extrapolating these patterns to distant kinematic space.

To help measure this effect, we use a standard 70%/20%/10% split to separate the kinematic bins into three groups, which we will call the training set, validation set and test set, respectively. The separation by kinematic bin (25 points in φ collected by x_{Bj} , t , Q^2 , E_b) is done randomly for each round of training. Note that this differs from the standard ML practice of splitting the data points themselves into random groups. The training set is used to directly adjust the parameters of the model. The validation set is used as an indicator of performance during

FIG. 1. Top: kinematic region in x_{Bj} and t and (bottom) x_{Bj} and Q^2 where experimental measurements used in this analysis have been taken at 6 and 12 GeV.

training—either for early stopping (see Sec. III C), or for hyperparameter selection.¹ Final performance will be reported on the test set, which is made up of kinematic bins with which the model has never been trained or evaluated. A model that performs well on the test set is capable of predicting entire φ curves in unseen regions of kinematic space, under the assumption that we are testing on inputs drawn from the same distribution as our training data. This assumption often breaks down in practice, as we will see in Sec. III D. Generalizing to unseen kinematic bins is a more difficult challenge than generalizing to held-out *data points* in bins that were part of the training set, because the overall shape of the φ curve can be quite predictable as long as nearby points can provide a good frame of reference.

Based on the analysis of [9], we experimented with transforming the dataset of $(x_{Bj}, t, Q^2, E_b, \varphi)$ by

¹A hyperparameter is a fixed constant that controls some aspect of an algorithm’s behavior. They can have a significant impact on performance, but cannot be learned like many of the model’s core parameters. As an example, SGD optimizes the parameters of a model, but the initial learning rate and batch size of the optimizer are predetermined hyperparameters.

concatenating $[\sin(\varphi), \cos(\varphi), \sin(2\varphi), \cos(2\varphi), \sin(3\varphi), \cos(3\varphi)]$. These harmonics could serve as additional features that assist the model in learning the cross section’s φ dependence. We will refer to the use of this additional input information as “harmonic features.”

During the training of the neural network the small value of the cross sections lead to numerical instabilities from floating-point underflows. The standard remedy for this obstacle is scaling the values during training; we scaled by 1000. Once training was completed we then return to the original scale for showing results.

B. Baselines

First, we examine the possibility of applying basic regression techniques to cross section fits to get an understanding of the difficulty of the problem and establish a performance baseline. We use the implementations provided in [17]:

- (i) *Linear* is a simple linear regression model. A vector of weights $\mathbf{W}$ and a scalar bias b map the input kinematics, x , to the cross section prediction with $\hat{\sigma} = \mathbf{W}x + b$. These parameters are the solution to a closed-form system of equations that minimizes the mean squared error (MSE) of the model’s predictions [18].
- (ii) *SVR* is a support vector machine regression model [19]. The model parameters, $\mathbf{W}$ and b , are the solution to the optimization problem [20]:

$$\begin{aligned} \min_{\mathbf{w}, b, \gamma, \gamma^*} \quad & \frac{1}{2} \mathbf{w}^T \mathbf{w} + C \sum_{i=1}^l \gamma_i + C \sum_{i=1}^l \gamma_i^* \\ \text{subject to} \quad & \mathbf{w}^T \rho(\mathbf{x}_i) + b - \sigma_i \leq \epsilon + \gamma_i, \\ & \sigma_i - \mathbf{w}^T \rho(\mathbf{x}_i) - b \leq \epsilon + \gamma_i^*, \\ & \gamma_i, \gamma_i^* \geq 0, i = 1, \dots, l, \end{aligned}$$

where C and ϵ are hyperparameters that control error tolerance and ρ is a kernel function that is used to map input vectors to a space where they may be more linearly separable. We experiment with both polynomial and radial basis function kernels.

1. Evaluating performance

Because of the large range of cross section targets present in our dataset, common regression metrics like mean squared error are biased towards a small number of kinematic bins. Therefore, we focus on statistics that normalize for the value of the true cross section. The mean absolute percentage error is a better comparison, but suffers from outliers when the model is inaccurate on very small cross sections. Instead, we compute the median absolute percentage error (MAPE). While this is an improvement, the dataset has enough noise that relatively accurate fits can have high error. With this in mind, we also compute an

accuracy metric, which we define to be the percentage of predictions that are close enough to the target value to be within the experiment's error bars defined according to Eqs. (8) and (9):

$$\text{Accuracy} = \frac{1}{N} \sum_i \Delta_{\sigma_i} \quad (8)$$

$$\Delta_{\sigma_i} = \begin{cases} 0, & \text{if } \begin{cases} \sigma_i - \hat{\sigma}_i > 0 & \text{and } |\sigma_i - \hat{\sigma}_i| > \delta_{\hat{\sigma}_i} + \delta_{\sigma_{i,\uparrow}} \\ \sigma_i - \hat{\sigma}_i < 0 & \text{and } |\sigma_i - \hat{\sigma}_i| > \delta_{\hat{\sigma}_i} + \delta_{\sigma_{i,\downarrow}} \end{cases} \\ 1, & \text{else } \begin{cases} \sigma_i - \hat{\sigma}_i > 0 & \text{and } |\sigma_i - \hat{\sigma}_i| < \delta_{\hat{\sigma}_i} + \delta_{\sigma_{i,\uparrow}} \\ \sigma_i - \hat{\sigma}_i < 0 & \text{and } |\sigma_i - \hat{\sigma}_i| < \delta_{\hat{\sigma}_i} + \delta_{\sigma_{i,\downarrow}} \end{cases} \end{cases} \quad (9)$$

We use an accuracy metric which is defined by the number of points that lie within experimental error bars. Δ_{σ_i} is assigned a value of +1 or 0 based on the difference between the cross section prediction, $\hat{\sigma}_i$, and the data, σ_i . The assignment is determined by a flowchart of criteria. First we determine whether the difference is positive or negative, indicating whether we are looking at the positive side of the error bar or the negative. Then we compare the absolute value of the difference $|\sigma_i - \hat{\sigma}_i|$ to the size of the experimental error, $\delta_{\hat{\sigma}_i}$, plus asymmetric systematic errors, $\delta_{\sigma_{i,\uparrow}}$, $\delta_{\sigma_{i,\downarrow}}$. If the absolute difference is less than the size of the error bar then we assign a +1 value to Δ_{σ_i} , else we assign a 0. The accuracy is then defined as the sum of all Δ_{σ_i} divided by total number of data points.

After a hyperparameter search, the best model is selected according to its MAPE on the validation set; it is then evaluated for the first time on the test set. The results are shown in Table II.

Overall, the baselines perform poorly. While the linear models' high error is hardly a surprise given the cross section's trigonometric shape across φ , the SVR model with access to high degree polynomial features still struggles to find an accurate fit. Results on the use of additional harmonic features are mixed, and we decide

against using them in the experiments that follow. In general, high error metrics like these highlight the complexity of the cross section function across the whole range of kinematic bins. Simple approaches are not sufficient for accurate predictions.

C. Training a deep network on DVCS data

We train two separate NNs, one to predict σ_{UU} and the other to predict σ_{LU} , we report the results of the training curves in Fig. 2. There are several major challenges in applying Deep Learning to this problem. The first is the range of the cross section data values; the target values for our networks' predictions span seven orders of magnitude in both the polarized and unpolarized cross sections. While the last layer of our architecture uses a linear activation function that can output any scalar, this wide range in targets destabilizes training. When kinematic points with large cross sections are sampled in a batch of training data, they can create outlier loss values that increase the size of our update gradients and make large adjustments to model parameters that decrease performance in regions with a smaller cross section. This effect is magnified by the MSE function, which is quadratic in the absolute error of our models' predictions. We address this by changing our loss function to the Huber loss:

$$\mathcal{L}_\theta^H(\mathbf{x}, \sigma) = \begin{cases} \frac{1}{2}(\sigma - f_\theta(\mathbf{x}))^2, & \text{if } |\sigma - f_\theta(\mathbf{x})| \leq \delta; \\ \frac{1}{2}|\sigma - f_\theta(\mathbf{x})|, & \text{if } |\sigma - f_\theta(\mathbf{x})| > \delta. \end{cases} \quad (10)$$

The Huber loss is quadratic in the model's error when that error is less than a hyperparameter δ , and linear otherwise. This limits the effect of large cross section values relative to MSE. The median value of the cross section data is relatively small, so the squared error term is stable for most data points and therefore the specific choice of δ is unimportant; we fix $\delta = 1 \text{ nb/GeV}^4$. Many data points have error bars that are listed well above 100% of the cross section value. Fitting to the center of such a wide range can create irregularities that generalize poorly and encourage overfitting. Ideally, our model would pay less attention to

TABLE II. Results of baseline methods on the test set. Each training run uses an identical Train/Val/Test split. We focus on median absolute percentage error because the large range of cross section sizes biases metrics like mean squared error towards a small number of kinematic bins. Accuracy is the % of predictions that lie within experimental error bars. Best results are in bold for clarity.

Method	UU				LU			
	Standard		Harmonic features		Standard		Harmonic features	
	Median % error	Accuracy (%)	Median % error	Accuracy (%)	Median % error	Accuracy (%)	Median % error	Accuracy (%)
Linear	238.87	2.10	343.68	1.0	293.03	19.68	333.23	22.34
SVR	45.73	19.14	37.57	27.13	68.09	57.71	72.97	57.98
Theory	13.39	49.50	...	...	62.46	62.67	...	...
<i>FemtoNet</i>	9.99	61.97	...	...	56.06	63.03	...	...

FIG. 2. Training curves for FemtoNet LU (top) and UU (bottom). The train loss is computed across batches of the training set while the validation loss is measured at the end of each epoch. When the validation loss has not hit a new minimum in the last 20 epochs, training ends to avoid overfitting. Notice that the validation outperforms the training due to the advantage of drop-out being turned off. The early stopping technique terminates training before overfitting.

highly uncertain measurements. Towards this goal, we consider a variation on the standard Huber Loss that incorporates the relative size of the experimental error bars for each data point. Let $\Delta\sigma$ represent the size of the symmetric error bar, while $\Delta\sigma_{\downarrow}$ and $\Delta\sigma_{\uparrow}$ represent the negative and positive asymmetric error, respectively. The error weighted Huber loss (EW Huber) is then defined:

$$\mathcal{L}_{\theta}^{\text{EW}}(\mathbf{x}, \sigma) = \frac{\mathcal{L}_{\theta}^H(\mathbf{x}, \sigma)}{1 + ((2\Delta\sigma + \Delta\sigma_{\downarrow} + \Delta\sigma_{\uparrow})/\sigma)}. \quad (11)$$

The EW Huber discounts the loss value of data points where the total size of the error bar is large relative to the scale of the cross section. Points with no experimental error stay unchanged and the effect of the weighting increases with the experimental uncertainty.

In many regions of kinematic space, the unpolarized cross section is dominated by the Bethe-Heitler component,

FIG. 3. Comparing several of our design choices across varying model sizes. “MSE” uses a standard mean squared error loss functions. “Huber” switches to a clipped variant of MSE that protects against outliers. “BH Ratio” generates cross sections by multiplying the output of the network by the value of the Bethe-Heitler component, which is calculated based on known theory [9] inside of a custom layer. “EW Huber” reweights terms in the Huber loss based on the size of the experimental error bars—reducing the importance of imprecise data points. Gradient clipping is used across all experiments.

which can be computed as a known function of the network’s inputs. We take advantage of this by implementing the Bethe-Heitler calculation as a custom neural network layer with fixed parameters that cannot be adjusted by gradient descent. This layer computes the Bethe-Heitler component of the input point, which is then multiplied by the output of the standard NN architecture to produce the predicted cross section. This effectively changes the function the NN must learn to approximate from σ to the simpler σ/BH . We test these approaches on increasingly large NN sizes to see how they affect performance, and the results are shown in Fig. 3. The Huber loss turns out to be essential for successful training. Performance improves with model size—perhaps the biggest advantage of deep learning techniques. The BH ratio simplifies the NN function enough to help small models, but as models grow larger their representational capacity improves and this simplification becomes unnecessary.

1. Overfitting and early stopping

Our models’ goal is to make accurate predictions in previously unseen kinematic bins. By improving their performance on the training set through gradient descent, we tend to improve their performance on the test set. However, over-optimizing on those training points can eventually come at the expense of test set performance. This effect is known as “overfitting” and is an important concern when training deep networks. A simple and effective heuristic to stop training before overfitting occurs is to periodically evaluate the model on the *validation* set.

FIG. 4. A schematic of the FemtoNet model architecture. Two large hidden layers with ReLU activations feed into a linear output layer that determines the cross section prediction. Dropout is applied after each fully connected (FC) layer.

When the validation loss stops decreasing, we are in danger of overfitting, and end the training run. This approach, known as early stopping, uses the validation set as a stand-in for the test set. SGD assumes the training, testing, and validation sets are sampled from the same underlying data distribution. Therefore, we can use the performance on the validation set as a sample of our performance on the testing set, as long as we are not making gradient updates on the validation data. We use early stopping in every training run, ending training after the validation loss has not decreased in more than 20 full iterations through the training set (epochs).

The size of the dataset is another obstacle. While many of the most successful applications of DL use training sets of millions [21] or even billions [22] of samples, the DVCS experiments at our disposal leave us with a little under 4,000 data points. Small datasets increase the risk of overfitting and often force us to *regularize* the model, or to restrict its ability to move freely in parameter space and arrive at a solution that is hyperspecific to the training set. We accomplish this with dropout regularization [23] implemented in our NN architecture as shown in Fig. 4. With dropout, connections between neurons in the network are severed before each forward pass with some probability p , which becomes an important hyperparameter. This is typically implemented by sampling binary weight matrices that serve as masks of the network weights—setting them to zero when the connection should be removed and one otherwise. This is shown pictorially in Fig. 5. Intuitively, Dropout reduces overfitting by preventing the network from being overly reliant on the outputs of a small number of neurons. There are many other approaches for regularization. We focus on dropout because of its use in predicting model uncertainty, which will be discussed further in Sec. IV.

2. FemtoNet

The final architectures of our cross section prediction networks are determined by a large search, guided by the Hyperband algorithm [24]. Hyperband trains many possible variations of the network architecture and uses the

FIG. 5. Example of the application of a binary dropout mask to a fully connected neural network. The 1s and 0s are assigned based on a probability distribution over each weight. When we resample the binary mask over each forward pass of the network, the variance of the output can be used as an approximation of statistical uncertainty of the neural network predictions.

results to decide which kinds of networks it should dedicate more time/resources to investigating and which kinds to disregard. The final UU and LU models both use the Adam optimizer [15] with learning rate 7.5×10^{-4} . The dropout probability $p = 0.25$. FemtoNet UU uses two hidden layers, both of size 2048, while FemtoNet LU uses two hidden layers with sizes of 2048 and 1024. Both networks use the EW Huber loss. They also normalize inputs before the first layer, with mean and standard statistics determined from the training set.

We evaluate FemtoNet on the same test set as the baseline approaches, with the results also listed in Table II. FemtoNet comfortably outperforms the ML baselines, though the improvement is much more significant in the unpolarized cross section than the polarized data. We discuss this finding further in Sec. IV C. In addition to comparing against other ML techniques, we compare against the theoretical cross section model developed in [9]. The NN is even more accurate over the test set—adding DVCS cross section prediction to a long list of problems in ML in which the performance of hand crafted or manually specified heuristics or formulas can be matched by *learned* features that are automatically discovered. However, the theory model has the important advantage of consistency across kinematics far outside the region of the training data, while the NN is less constrained.

In the upper plot of Fig. 6, we compare the cross section predictions of FemtoNet to the theoretical predictions across a slice of t and ξ . The region of existing experimental data is indicated by the shaded rectangle. Unsurprisingly, the two agree quite strongly in this area, where both have already been proven to be accurate based on the test set evaluation. However, as we move away from existing data we see a dramatic spike in the difference in the NN's predictions relative to the theory-based model. Note that this experiment says nothing about the accuracy of the theory's predictions over these kinematics—there is no experimental data here to serve as the ground truth. All we can conclude is that the theory and ML significantly

FIG. 6. Agreement between the theory-based cross section model and the NN approximation before (top) and after (bottom) the addition of generated pseudodata. This is calculated fixing $Q^2 = 2 \text{ GeV}^2$ plotted as a function of the variables momentum transfer variables ξ and t . The shaded rectangle represents where current DVCS exclusive data sit at a $Q^2 \approx 2 \text{ GeV}^2$.

disagree. The ability of NNs to generalize outside their training distributions is an active area of research in ML. Deep learning is often concerned with high dimensional data, such as images or audio files, where any concept of distance between data points is essentially meaningless. In contrast, the x_{Bj} , t , and Q^2 variables of kinematic space make it easy to think about “how far” our model can generalize, and may make this problem of interest to the broader ML community. Figure 6 suggests that FemtoNet learns a generalizable representation of the t dependence, but struggles with other kinematic variables.

D. Scaling to larger datasets

While the cross section measurements used in our analysis consist of around 4,000 data points, that total will

likely see significant growth from future experiments. We are interested in developing a method that can scale as new measurements are collected. This can be tested by expanding the kinematic region to include a larger range of x_{Bj} , t , and Q^2 . In particular, we generate pseudodata in regions outside the range where data currently exists to measure the network’s compatibility with theoretical calculations of the cross section.

1. Pseudodata

In Fig. 7 we demonstrate the process for generating pseudodata using several diquark model calculations to parametrize the GPDs. The panels are described in descending order in detail in the text from subsections (a)–(d).

(a) GPD envelopes are generated using two parametric forms following [25] as H_{VA} , E_{VA} , $\tilde{H}_{VA}$, $\tilde{E}_{VA}$ and [26] as H_{VGG} , E_{VGG} , $\tilde{H}_{VGG}$, $\tilde{E}_{VGG}$. To ensure an unbiased and suitably “random” choice of the parametric form for the GPD we vary the parameters P_i^o of each model by a random fluctuation, Δ_i , of up to $\pm 10\%$ of P_i^o dictated by a random number generator. We create the parameters of the model in the following way:

$$P_i = P_i^o + \Delta_i. \quad (12)$$

This is done for both parametric forms of the GPDs. The final envelope is generated by combining the two models in a random combination,

$$F = \alpha F_{VA} + \beta F_{VGG} \quad (13)$$

where $F = [H, E, \tilde{H}, \tilde{E}]$, α is a randomly generated number between 0 and 1, and $\beta = 1 - \alpha$. Using this method, for each kinematic bin we generate 25 GPD curves as shown in the top panel of Fig. 7.

(b) We evolve the GPDs using pQCD evolution equations and convolute the GPDs with the corresponding Wilson coefficient function to calculate an envelope of 25 real and imaginary parts of the Compton form factors. The imaginary part of the Compton form factor is directly related to the value of the GPD at the crossover point between Efremov-Radyushkin-Brodsky-Lepage (ERBL) and Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) regions multiplied by π , therefore it is more susceptible to the random fluctuations of the GPDs. This is in contrast to the real part of the Compton form factor which is constrained by the enveloping Wilson coefficient function and integral over the longitudinal momentum variable x .

(c) The CFF envelope enters into the cross section following the theory discussed in Sec. II above and [9,27]. We generate pseudodata by evaluating the cross section at azimuthal angle, φ , values starting at 7.5 and going in increments of 15 degrees until we reach 360. For each of the 24φ values we generate 25 cross section points using each of the 25 generated CFFs. This creates an

FIG. 7. Process for creating pseudodata at the kinematics $Q^2 = 1.820 \text{ GeV}^2$, $x_{Bj} = 0.343$, $t = -0.172 \text{ GeV}^2$, $E_b = 5.75 \text{ GeV}$: (a) construction of GPD envelopes, (b) generation of CFFs from GPD envelopes, (c) creation of an envelope of cross section values, (d) final pseudodata with simulated errors.

envelope of cross section calculations for each kinematic bin. The spread of the cross section values can be seen to be suitable enough to be random, but it is approximately centered around the actual data values.

(d) For each φ value of the data, we randomly pick one data point from the 25 in the cross section envelope. The random choosing of the data also simulates some of the statistical fluctuations naturally present in the experimental data. We generate the pseudodata errors by comparing to the dataset in the regions where there is data and controlling the error to match the relative size of the cross section errors. Using pseudodata, we have control over the size of the error bars giving us a unique opportunity to study the effects of experimental error when using a neural network analysis. We then use this metric to extend into regions where there is no data, simulating approximate errors in regions unexplored by experiment. The coverage of our pseudo-dataset relative to existing experimental data is shown in Fig. 8.

2. Evaluating performance on the pseudo-dataset

We evaluate our deep learning approach on an expanded pseudo-dataset of more than 30,000 unpolarized cross

FIG. 8. Kinematic regions in (top) x_{Bj} versus Q^2 and (bottom) x_{Bj} versus t where we generate pseudodata compared to where current DVCS data exists.

TABLE III. Pseudo-dataset results. NN approaches scale well to a dataset of more than 30,000 samples.

	Median % error	Accuracy (%)
Linear	3090.56	1.50
SVR	42.46	47.75
<i>FemtoNet</i>	12.50	73.90

section measurements. One of the most important advantages of NNs is their scalability; it is not uncommon for hard problems in ML to be solved by simply scaling up a model’s parameter count. This is certainly an option when fitting cross sections. However, for the sake of comparison, we use the same training process and model architecture developed for the experimental datasets. The results are shown in Table III. Our NN pipeline adapts well, highlighting its ability to learn an accurate approximation of the cross section across a large region of kinematic space where it is more difficult to rely on overfitting or “memorization.”

In the bottom of Fig. 6, we investigate the pseudodata’s impact on the shape of the NN’s approximation relative to the theory-based model described in [9]. As discussed in Sec. III C, the NN can diverge significantly from the theoretical model when predicting on kinematic bins that are far from existing measurements. When given access to cross sections in more distant regions of kinematic space (as part of the pseudo-dataset), this problem vanishes, underscoring the importance of collecting DVCS measurements across a wide range of kinematics in future experiments.

IV. PREDICTING THE CROSS SECTION

A. Estimating model uncertainty during inference

When making predictions that have an effect on decisions or later calculations, we would like to know the confidence level of our model—how sure are we that our predictions are accurate? This confidence is quantified as model uncertainty. Recent literature (see Refs. [28,29] for DVCS, [30,31] and references therein for inclusive scattering) has adopted the “replica method” as a means of estimating uncertainty. The replica method is similar to the bootstrap aggregating [32] approach that is common in ML, in which a set of models is trained by sampling from the dataset with replacement and using the disagreement of their predictions to estimate uncertainty. The replica method consists of running the learning algorithm many times on versions of the dataset where points are randomly perturbed relative to their experimental error bars. This is meant to propagate the experimental error into the results of the model. After many cycles of resampling and retraining, we have a set of models trained on slightly different datasets. The mean output of this set of models is used as the prediction, and the standard deviation serves as an approximation of uncertainty. While this is a valid way to

incorporate experimental/epistemic uncertainty into the modeling process, the repeated training can be prohibitively expensive when working with large models or datasets. Using multiple copies of the model parameters can be memory intensive, difficult to vectorize, and less computationally efficient when making predictions. Ideally, we could recover model uncertainty from a single neural network.

One possibility is a Bayesian neural network, where we optimize a distribution over weights and sample from that distribution between each prediction [33]. While effective, these networks have significantly more parameters and can be more difficult to train. A simple alternative is to leave dropout turned on during inference. By resampling our binary mask over the weights between each forward pass, we create a prediction distribution that can be used to approximate model uncertainty. This technique can be shown [34] to be an approximation of a Gaussian process. Our model generates a batch of k predictions over the same data point, each with different dropout masks. Unless otherwise noted, we set $k = 1000$. This smooths the uncertainty bands for plotting, but this value can be decreased if inference speed is a priority.

B. Learning the ϕ dependence

First, we examine the model’s ability to fit the range of ϕ values in existing kinematic bins (x_{Bj} , t , Q^2 , E_b). FemtoNet’s predictions in unseen kinematic bins from the test set are shown in Fig. 9. The models learn to predict the cross section within experimental error bars, and are robust to noise and outliers, particularly in the Hall B datasets.

C. Unpolarized vs polarized results

The total unpolarized cross section at leading twist contains three pieces as shown in Eq. (2): a piece that is bilinear in Compton form factors (DVCS), an interference piece that is linear in Compton form factors ($\mathcal{I}$), and a piece that has no dependence on Compton form factors and is an exactly calculated background parametrized by the elastic form factors (BH). The Bethe-Heitler cross section term, in most kinematic regimes, dominates the total cross section; often an order of magnitude or larger than the interference and DVCS terms combined. We can calculate the background to a well-known degree of accuracy using covariant four vector products. The error and the noise of the cross section is thus due entirely to the terms which are parametrized by the Compton form factors. Since the background dominates the cross section, it then also dominates the noise inherent in scattering cross section data, effectively reducing the errors and simplifying the functional form. It is demonstrated in Table II that the neural network can predict the unpolarized cross section with a relatively small median percentage discrepancy.

FIG. 9. The ϕ dependence of σ_{LU} (top row), and σ_{UU} (bottom row) in previously unseen kinematic bins. Shading indicates two standard deviations of 1000 predictions.

This is in contrast to the polarized cross section which has a drastically different result. The functional form of the polarized cross section we find is more difficult for the neural network to learn. We can explain this by the presence of a large contribution of noise in the polarized data. By parity constraints, the longitudinally polarized beam on an unpolarized target (LU) cross section has no Bethe-Heitler background nor a term bilinear in Compton form factors at leading twist. It is linear in Compton form factors; as such, there is no error/noise suppression as in the unpolarized case.

We can compare the LU polarized cross section to an equivalent reduced unpolarized cross section where we have subtracted off the Bethe-Heitler background to investigate the impact of the dominant term on the errors as shown in Fig. 10. Notice that the unpolarized cross section tends to converge to a relative error similar to that of the polarized cross section at similar kinematics, where the relative error is defined as

$$\bar{\delta}_\sigma = \frac{1}{N} \sum_i \frac{\delta_{\sigma_i}}{\sigma_i}. \quad (14)$$

The background in the unpolarized cross section tends to simplify the functional form of the cross section while simultaneously smoothing out the noise and decreasing relative error. This, however, presents a problem for extracting Compton form factors where the cross section terms they parametrize tend to have large noise components.

D. Predictions in new kinematic regions

We are not restricted to kinematic bins that have already been measured—the model learns to extrapolate between existing bins. We can sweep through regions where data does not exist, and use the model’s uncertainty estimate to advise our use of its predictions. In Fig. 11 we illustrate this point for two particular kinematic values.

FIG. 10. Top: the total polarized cross section which at leading twist is linearly dependent on Compton form factors. Bottom: the reduced unpolarized cross section which we define as the total cross section subtracting the Bethe-Heitler cross section. The reduced cross section has two terms both dependent on Compton form factors. Notice that the relative errors of these two converge without the presence of a background process in the unpolarized case.

This result is consistent with what was reported in the upper panel of Fig. 6 where the model’s uncertainty in the experimental data domain is very low. Notice that the neural network seems to be able to predict the t dependence much more accurately than the ξ dependence. This result could mean, on one side, that the t dependence has a simpler and consistent functional form throughout the phase space than the ξ dependence. The latter will, in fact, be affected by perturbative QCD evolution. On the other hand, the kinematic spread of the data for the t variable covers a wider section than for ξ ; this feature might also influence the outcome. Clearly, the availability of more data in a wider range in ξ for fixed t will help in interpreting this result.

While the dataset contains thousands of data points, it contains significantly fewer unique kinematic bins; most of the individual points come from kinematic sweeps in the azimuthal angle ϕ . The impact of this can be seen empirically in the model’s uncertainty estimates. The uncertainty grows quickly as we move between different kinematic regions, while it remained relatively constant between different ϕ values in a fixed region.

FIG. 11. Predictions across a range of ξ values. ξ is converted to x_{Bj} before being sent to the model. The blue dots in the figure represent the experimental data. The uncertainty in the curves was discussed in Sec. IV A. Notice that for $\xi < 0.1$ the cross section is evaluated in a kinematically unphysical region where the neural network can still make predictions.

V. CONCLUSIONS AND OUTLOOK

In summary, a deep learning neural network analysis is essential to extract information from deeply virtual exclusive scattering data in a model independent way. These reactions are characterized by final states with a large number of kinematic variables and observables, and a consequent growth in complexity of the mapping between the initial variables into the output states.

Our analysis is a necessary first step in a future global extraction of Compton form factors from data. The extraction of information from the data requires an understanding of the features that directly impact such an analysis. In our case, the significant contribution of noise and the relatively small kinematics spread of the data made predicting the cross section a difficult problem even using a deep neural network. Nevertheless, as shown in Table II, an analysis using a DNN was proven the most efficient and accurate method, as compared to other baselines, for extracting

information. Several additional challenges in our analysis stem from the current lack of an extensive centralized experimental database for exclusive scattering data. Another difficulty with this dataset is the mismatch in the quantity of data between the φ dependence and the rest of the kinematic variables (t , Q^2 , and x_{Bj}). This results in the neural network not generalizing well across kinematic space. Such challenges can be curtailed through the establishment of a common dataset with benchmarking requirements for the publishing of data. This would optimize the usage of data from different experiments in global analyses.

Our results from FemtoNet suggest many interesting characteristics of DVCS data. In particular, we have shown that there is a substantial and consistent signal in the unpolarized cross section. This is at variance with the polarized cross section where the median absolute percentage error is much larger (see Table II). Furthermore, our results extrapolate to kinematics outside of the range covered in experiment: the DNN seems to be able to generalize in the kinematic variable t more precisely and for a wider range than for the other variables ξ and Q^2 . This suggests, one side, interesting implications for Fourier transforms for 3D visualization from DVCS data. On the other hand, we would like to express a word of caution concerning the DNN's ability to generalize beyond existing data. The apparent difficulty in predicting the ξ dependence could, for instance, emerge from the DNN's treatment of the phase space cutoffs in the data. Further insight on the behavior of the DNN in predicting the cross section's trends could be obtained with future, more extensive datasets.

We also studied the question of how to use experimental errors in neural networks by incorporating them in the error weighted Huber loss [EW Huber, Eq. (11)], and also in our accuracy metric (Eqs. (8) and (9)). For this purpose we used pseudodata. With pseudodata we have control over the error bars and can establish relationships between

experimental errors and the uncertainty on the extracted Compton form factors.

In future studies we will, on one side, extend our approach to analyze other exclusive experiments at Jlab@12 GeV, as well as at the Electron Ion Collider, while also addressing remaining open questions. We plan to further study both the accuracy and the uncertainty obtained from the DNN analysis of DVCS data. In particular, we will analyze the contribution of the background Bethe-Heitler process in the unpolarized cross section that seems to dominate the accuracy of the predictions with respect to the polarized electron scattering. The framework provided in this paper to treat the experimental error in ML affords us the ability to next study the effect of the data errors on the uncertainty of the extracted Compton form factors. This will, in turn, allow us to optimize both the precision and the kinematic domain for experimental measurements at future colliders needed to efficiently disentangle the observables from the data.

Finally, our method is flexible to be extended to incorporate lattice QCD results in the data analysis. The latter will play an even more important role in future exclusive experiment analyses than in deep inelastic scattering, towards the ultimate extraction of GPDs from data.

ACKNOWLEDGMENTS

This work was funded by DOE Grant No. DE-SC0016286 and in part by the DOE Topical Collaboration on TMDs (B.K. and S.L.), DOE Grant No. DE-FG03-95ER40965 (M.B.) and SURA Grants No. C2019-FEMT-002-04 and No. C2020-FEMT-006-05. We gratefully acknowledge the University of Virginia School of Data Science for logistical support; technical assistance by Bryan Wright and Rick Marshall (UVA, Physics), and many insightful physics discussions with Gordon Cates.

-
- [1] A. Ohnishi, Y. Mori, and K. Kashiwa, J. Phys. Soc. Jpn. Conf. Proc. **26**, 024011 (2019).
 - [2] Y. Mori, K. Kashiwa, and A. Ohnishi, Phys. Rev. D **96**, 111501 (2017).
 - [3] A. Alexandru, P. F. Bedaque, H. Lamm, and S. Lawrence, Phys. Rev. D **97**, 094510 (2018).
 - [4] P. Bedaque *et al.*, arXiv:2006.05422.
 - [5] K. Goeke, M. V. Polyakov, and M. Vanderhaeghen, Prog. Part. Nucl. Phys. **47**, 401 (2001).
 - [6] M. Diehl, Eur. Phys. J. C **19**, 485 (2001).
 - [7] A. Belitsky and A. Radyushkin, Phys. Rep. **418**, 1 (2005).
 - [8] K. Kumericki, S. Liuti, and H. Moutarde, Eur. Phys. J. A **52**, 157 (2016).
 - [9] B. Kriesten, S. Liuti, L. Calero-Diaz, D. Keller, A. Meyer, G. R. Goldstein, and J. Osvaldo Gonzalez-Hernandez, Phys. Rev. D **101**, 054021 (2020).
 - [10] FemtoNet, Neural networks for nuclear femtography, https://github.com/uva-femtography/FemtoNet/tree/master/dvcs_cross_section (2020).
 - [11] H.-S. Jo, Proc. Sci., QNP2012 (2012) 052 [arXiv:1207.3709].
 - [12] M. Defurne *et al.* (Jefferson Lab Hall A Collaboration), Phys. Rev. C **92**, 055202 (2015).
 - [13] M. Defurne *et al.*, Nat. Commun. **8**, 1408 (2017).
 - [14] F. Georges, Deeply virtual Compton scattering at Jefferson Lab, Ph.D. thesis, Institut de Physique Nuclaire d'Orsay, France, 2018.

- [15] D. P. Kingma and J. Ba, [arXiv:1412.6980](#).
- [16] G. Hinton, N. Srivastava, and K. Swersky, Neural networks for machine learning, Lecture 6a: Overview of mini-batch gradient descent (2012), <http://www.cs.toronto.edu/~hinton/coursera/lecture6/lec6.pdf>.
- [17] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *J. Mach. Learn. Res.* **12**, 2825 (2011).
- [18] J. F. Kenney and E. Keeping, *Mathematische Operationsforschung und Statistik* **1**, 252 (1962).
- [19] V. N. Vapnik, *IEEE Trans. Neural Networks* **10**, 988 (1999).
- [20] C.-C. Chang and C.-J. Lin, *ACM Trans. Intell. Syst. Technol.* **2**, 1 (2011).
- [21] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, in *2009 IEEE conference on computer vision and pattern recognition* (IEEE, 2009), pp. 248–255.
- [22] T. B. Brown *et al.*, [arXiv:2005.14165](#).
- [23] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [24] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, *J. Mach. Learn. Res.* **18**, 1 (2018).
- [25] G. R. Goldstein, J. O. G. Hernandez, and S. Liuti, *J. Phys. G* **39**, 115001 (2012).
- [26] M. Vanderhaeghen, P. A. M. Guichon, and M. Guidal, *Phys. Rev. D* **60**, 094017 (1999).
- [27] B. Kriesten and S. Liuti, [arXiv:2004.08890](#).
- [28] K. Kumericki, D. Mueller, and A. Schafer, [arXiv:1110.3798](#).
- [29] H. Moutarde, P. Sznajder, and J. Wagner, *Eur. Phys. J. C* **79**, 614 (2019).
- [30] R. D. Ball *et al.* (NNPDF Collaboration), *Eur. Phys. J. C* **77**, 663 (2017).
- [31] R. D. Ball, V. Bertone, F. Cerutti, L. Del Debbio, S. Forte, A. Guffanti, J. I. Latorre, J. Rojo, and M. Ubiali (NNPDF Collaboration), *Nucl. Phys.* **B855**, 153 (2012).
- [32] L. Breiman, *Mach. Learn.* **24**, 123 (1996).
- [33] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, [arXiv:1505.05424](#).
- [34] Y. Gal and Z. Ghahramani, [arXiv:1506.02142](#).