

Parametrizing the detector response with neural networks

S. Cheong,^{a,b} A. Cukierman,^{a,b} B. Nachman,^c M. Safdari^{a,b,1} and A. Schwartzman^b

^a*Elementary Particle Physics Division, SLAC National Accelerator Laboratory,
2575 Sand Hill Rd, Menlo Park, CA, U.S.A.*

^b*Physics Department, Stanford University,
450 Jane Stanford Way, Stanford, CA, U.S.A.*

^c*Physics Division, Lawrence Berkeley National Laboratory,
1 Cyclotron Rd, Berkeley, CA, U.S.A.*

E-mail: murtazas@stanford.edu

ABSTRACT: In high energy physics, characterizing the response of a detector to radiation is one of the most important and fundamental experimental tasks. In many cases, this task is accomplished by parameterizing summary statistics of the full detector response probability density. The parameterized detector response can then be used for calibration as well as for directly improving physics analysis sensitivity. This paper discusses how to parameterize summary statistics of the detector response using neural networks. In particular, neural networks are powerful tools for incorporating multidimensional data and the loss function used during training determines which summary statistic is learned. One common summary statistic that has not been combined with deep learning (as far as the authors are aware) is the mode. A neural network-based approach to mode learning is proposed and empirically demonstrated in the context of high energy jet calibrations. Altogether, the neural network-based toolkit for detector response parameterization can enhance the utility of data collected at high energy physics experiments and beyond.

KEYWORDS: Analysis and statistical methods; Data processing methods; Large detector-systems performance; Pattern recognition, cluster finding, calibration and fitting methods

ARXIV EPRINT: [1910.03773](https://arxiv.org/abs/1910.03773)

¹Corresponding author.

Contents

1	Introduction	1
2	Mathematical overview of training neural networks	2
3	Learning different properties of the response distribution	3
3.1	Mean	3
3.2	Median	3
3.3	Mode	4
3.4	Standard deviation	6
3.5	Quantile	6
4	Numerical examples	7
4.1	Calorimeter and jet energy response	7
4.2	Muon momentum resolution	9
5	Discussion	10
6	Conclusions	11

1 Introduction

Any physics experiment is inevitably influenced by one’s ability to detect and measure physical observables of interest. In particular, an accurate understanding of the characteristics of detectors is important for designing an experiment, projecting the reach of an experiment, calibrating measurements, and ultimately analyzing the data collected. Therefore, one of the most fundamental and important experimental tasks is to understand the behavior of the detection apparatus. Beyond the physical energy deposited by radiation, measurements from real physics experiments are influenced by the electronic read-out processes as well as a variety of other hardware-, firmware-, and software-based algorithms. All of these effects are collectively called *the effect of a detector*.

Mathematically, any experimental measurement can be considered as a function that intakes some random variable X and outputs another random variable Y . The input is the ‘true’ quantity and the output is the ‘measured’ quantity. Due to non-negligible noise,¹ $Y|X$ is also a random variable and its distribution can be described by the probability distribution $p(y|x)$, where x and y are realizations of X and Y . Because the detector response is captured completely by the distribution $p(y|x)$, this distribution is referred to as *the response*.

For example, X could be the true energy of a particle traversing the detector, while Y is the reconstructed energy measured by the detector. X and Y need not have the same physical dimensions; for instance, a tracking detector immersed in a magnetic field aims to determine the momentum X of a charged particle by making spatial location measurements Y and calculating the sagitta of the particle trajectory. It is also worth noting that both X and Y are often multi-dimensional; that

¹Some of this noise is from marginalizing out unmeasured quantities while other noise is quantum-mechanical in origin.

is, a detector may measure multiple observables simultaneously and its response may depend on multiple variables including ones other than observables of primary interest.

For calibration and other experimental tasks, it is often necessary to parametrize detector effects with a function $\nu(x)$ that categorizes the response of the detector to an input x . The function ν encodes some summary statistic(s) of the probability distribution $p(y|x)$ and often represents the central tendency of the detector response such as the mean or the spread of the response such as the standard deviation. The parameterization $\nu(x)$ is typically obtained using detailed Monte Carlo (MC) simulations of an experiment or calibration data where the values of the input variable x are measured precisely by independent techniques.

A common method to obtain $\nu(x)$ is to bin the (simulated) data in X and then calculate ν based on a histogram of $Y|X$. This method provides a “look-up” table for values of $\nu(x)$ at the center of each bin of X , which can then be interpolated. While this method works well when X is low dimensional, it becomes infeasible when more observables are added to improve the accuracy of the detector description. In order to maintain the same statistical precision in a given X bin, the bin sizes would need to grow exponentially large for a linear increase in the dimensionality of X . Furthermore, interpolation between bins is more difficult in higher dimensions. The task is additionally challenging if Y is multidimensional.

An alternative to the bin method is to perform an unbinned fit using a functional form for $p(y|x)$ and $\nu(x)$. Such functional forms are nearly always an approximation and often do not capture all of the salient features of the data. For example, it is common to assume that $p(y|x)$ has a Gaussian core even though the tails are known to not be simply Gaussian. If the assumptions about $p(y|x)$ or $\nu(x)$ are inaccurate, any statistical fit will result in a biased parametrization of the detector, which will influence all consequent experimental studies. Hybrid methods using a functional form for $p(y|x)$ and a non-parametric form for $\nu(x)$ are also possible (see e.g. [1, 2]).

This paper describes methods for parameterizing the detector response using neural networks (NN’s). By their design, NN’s can readily incorporate unbinned high-dimensional inputs and outputs. The form of $\nu(x)$ will dictate the loss function used in the NN training. This connection will be explained in detail for the mean, median, standard deviation, and quantiles of $p(y|x)$. A new method (as far as the authors are aware) is proposed for a NN parametrization when $\nu(x)$ is the mode of $p(y|x)$.

This paper is structured as follows. A mathematical overview of NN training is provided in section 2. Then, section 3 discusses how to learn the various properties and obtain a parametrization $\nu(x)$ of the detector’s response using NN’s. In section 4, the proposed method is numerically demonstrated using realistic examples of a physical detector. Section 5 discusses the potential uses of the proposed method across experimental physics and a further application to neural-network-based calibration of detectors. The paper ends with conclusions and future outlook in section 6.

2 Mathematical overview of training neural networks

Neural network are trained by minimizing a loss function ℓ :

$$\nu = \operatorname{argmin}_{\nu'} \mathbb{E}[\ell(Y, \nu'(X))] \quad (2.1)$$

$$= \operatorname{argmin}_{\nu'} \iint \ell(y, \nu'(x)) p(x, y) dx dy, \quad (2.2)$$

where $\mathbb{E}[\cdot]$ represents the expected value or average. The form of ℓ controls what statistic of $p(y|x)$ will be parameterized by ν . One of the most common loss functions is the squared error $\ell(\nu(x), y) = \|y - \nu(x)\|_2^2$. Loss functions must be bounded below and typically have a unique local minimum. As it is the most common case and significantly simplifies the notation, the following discussion will focus on the case where Y is one-dimensional.

If the training dataset consists of a large enough unbiased sample from the distribution $p(x, y)$, then a sufficiently flexible NN and training procedure will arrive at eq. (2.1). Consider a small neighborhood $V(x_0)$ around some x_0 within the support of the training set. Assuming that the distribution $p(x, y)$ and the value of the NN function $\nu_0 \equiv \nu(x_0)$ are roughly constant (vary slowly) within $V(x_0)$, a successfully trained NN must satisfy:

$$\begin{aligned}
0 &= \frac{\partial}{\partial \nu_0} \mathbb{E}[\ell(Y, \nu(X))] \\
&= \frac{\partial}{\partial \nu_0} \left[\int_{V(x_0)} p(x) \int_{-\infty}^{\infty} \ell(y, \nu(x)) p(y|x) dy dx \right] \\
&\approx \left[\int_{V(x_0)} p(x) dx \right] \times \left[\int_{-\infty}^{\infty} \frac{\partial \ell(y, \nu_0)}{\partial \nu_0} p(y|x_0) dy \right] \\
&\propto \mathbb{E} \left[\left. \frac{\partial \ell(y, \nu_0)}{\partial \nu_0} \right| x_0 \right].
\end{aligned} \tag{2.3}$$

As ℓ is bounded from below, eq. (2.3) must be true everywhere across the support of $p(x)$ for a properly trained NN. Therefore, $\mathbb{E}[\partial \ell(y, \nu(x))/\partial \nu(x)|x] = 0$ which suggests that one can therefore engineer ℓ to target a statistic of interest for the response parameterization.

3 Learning different properties of the response distribution

The sections below discuss loss functions that can be used to learn some common statistics for response parametrization using NN's. The specific choice depends on the particular experimental application. Hybrid methods are also available, such as the Huber loss [3] which interpolates between the mean squared error and the mean absolute error. The asymptotic behavior of such approaches are more complicated than the ones presented in this section, but they may be useful in practice for certain situations.

3.1 Mean

The mean of the response distribution $p(y|x)$ can be learned using the mean square-error (MSE) loss: $\ell_{SE}(y, \nu(x)) \equiv (y - \nu(x))^2$. According to eq. (2.3), a NN trained successfully with ℓ_{SE} will satisfy:

$$\mathbb{E} \left[\left. \frac{\partial \ell_{SE}(y, \nu(x))}{\partial \nu(x)} \right| x \right] = 2(\nu(x) - \mathbb{E}[y|x]) = 0.$$

3.2 Median

The median may be a useful alternative to the mode for characterizing the central tendency of $Y|X$ because it is less sensitive to outliers. The median of the response can be learned using the mean

absolute-error (MAE) loss: $\ell_{\text{AE}}(y, \nu(x)) \equiv |y - \nu(x)|$. A NN trained with MAE will output $\nu(x)$ such that:

$$\mathbb{E} \left[\left. \frac{\partial \ell_{\text{AE}}(y, \nu(x))}{\partial \nu(x)} \right| x \right] = \mathbb{E} [\text{sgn}(y - \nu(x)) | x] = 0, \quad (3.1)$$

where $\text{sgn}(z)$ is the signum function defined by:

$$\text{sgn}(z) \equiv \begin{cases} 1 & \text{if } z > 0 \\ 0 & \text{if } z = 0. \\ -1 & \text{if } z < 0 \end{cases}$$

Then, eq. (3.1) can be more explicitly written as:

$$- \int_{-\infty}^{\nu(x)} p(y|x) dy + \int_{\nu(x)}^{\infty} p(y|x) dy = 0,$$

which precisely requires that $\nu(x)$ be the median of the distribution $p(y|x)$.

3.3 Mode

Like the median, the mode is insensitive to outliers. The mode has the additional property that it is insensitive to truncating the distribution on the left or right side, as is often necessary for realistic detector readout systems. However, unlike the median, the mode may not be unique. The mode of the response can be learned with $\ell(y, \nu(x)) = -\delta(y - f(x))$, where δ is the Dirac delta function. The solution to eq. (2.3) is shown using the properties of the Delta-function:

$$- \int_{-\infty}^{\infty} \frac{\partial \delta(y - \nu(x))}{\partial \nu(x)} p(y|x) dy = - \left. \frac{\partial p(y|x)}{\partial y} \right|_{y=\nu(x)} = 0.$$

The Dirac delta-function is an impractical loss function and so one must use a surrogate loss with similar properties. Various surrogates have been studied in the context of linear regression [4, 5], but as far as the authors are aware, this has never been combined with neural networks.

One possible surrogate to the Dirac delta-function is a negative Gaussian kernel:

$$\ell_{\text{GK}}(y, \nu(x); h) \equiv -\frac{1}{\sqrt{2\pi}h} \exp\left(-\frac{(y - \nu(x))^2}{2h^2}\right)$$

where h is a hyper-parameter called the ‘‘bandwidth’’ of the kernel.

Using a (negative) kernel function as a loss function to learn the mode is analogous to the core idea of kernel density estimation (KDE). KDE estimates a probability distribution by finite sum of kernels centered at a set of data points. Analogously, the mean training loss in a small neighborhood of given x using ℓ_{GK} will reproduce a smeared version of $-p(y|x)$ up to a normalization factor and a constant-offset. Therefore, the mean training loss is minimized when $\nu(x) = \text{argmax}_y p(y|x)$ across the support of the training dataset.

The potential of ℓ_{GK} is illustrated in figure 1, which shows the loss (analytically as in eq. (2.1)) landscape for a fixed $x = x_0$.

Plotted in figure 1 are the true mean losses for a given distribution of y :

$$\mathbb{E}[\ell(Y, \nu_0) | X = x_0] = \int_{-\infty}^{\infty} \ell(y, \nu_0) p(y|x_0) dy, \quad (3.2)$$

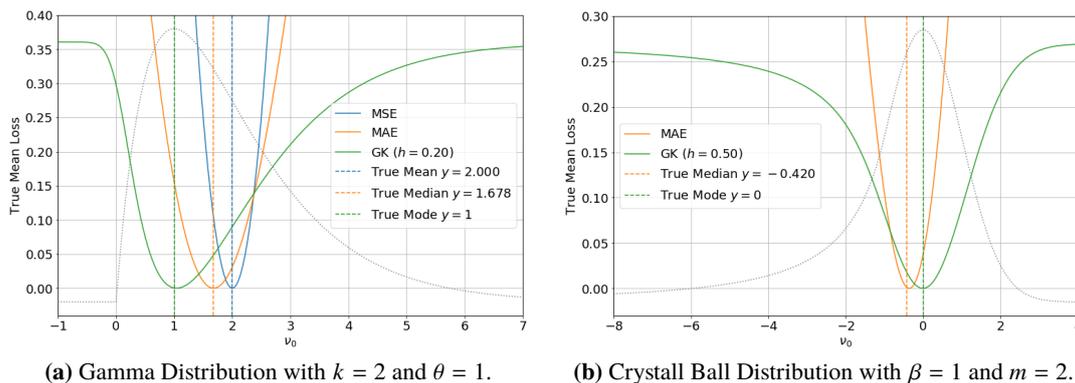


Figure 1. The true mean losses as a function of ν_0 for two different underlying distributions, each with a unique mode. Shown in gray dots are the underlying probability distributions, with arbitrary vertical scale. Constant off-sets have been subtracted from the loss functions such that their minima are all zero (purely for plotting purposes).

where $\nu_0 = \nu(x_0)$. For both the Gamma and Crystal Ball distributions, the mean GK loss mimics a smoothed version of the underlying probability density distribution. The GK loss does not reproduce the strict cut-off of the Gamma distribution at $y = 0$, but its minimum is still unbiased with respect to the true mode. As long as $p(y|x)$ is smooth around the mode, the function minimized with the GK loss predicts the mode. As advertised in section 3.1 and 3.2, figure 1 shows that the minima of MSE and MAE losses are at the mean and the median of $p(y|x)$, respectively. For such asymmetric response distributions, the mode is a good parameter with which to characterize the typical response of a detector, especially since the mean does not exist for the Crystal Ball distribution.

The negative Gaussian kernel may also not be a practical loss because its gradients vanish away from the mode, reducing the effectiveness of gradient-based optimization methods. One solution to this challenge is to introduce “leakiness” to the negative Gaussian kernel by defining the leaky Gaussian kernel (LGK) loss:

$$\ell_{\text{LGK}}(y, \nu(x); h, \alpha) = \ell_{\text{GK}}(y, \nu(x); h) + \alpha |y - \nu(x)|, \quad (3.3)$$

where α is a hyper-parameter called the leakiness. Note that the loss is still minimized (zero) when $\nu(x) = y$. The leakiness should be $\alpha \ll 1$ so that the value of $\nu(x)$ that solves eq. (2.3) is still close to the mode of $p(y|x)$. This term ensures a lower bound on the magnitude of the loss gradient, even when the NN output is far away from the optimal target value.

Figure 2 shows the true mean losses (eq. (3.2)) with various values of α for the gamma distribution with $k = 2$ and $\theta = 1$. Note that the mean LGK loss is minimized at values very close to the true mode (which is at $y = 1$) for small values of α . However, when α is large, the LGK loss approaches the absolute error $\ell_{\text{AE}}(y_1, y_2) \equiv |y_1 - y_2|$, and the minimum therefore gets biased towards the median at $y \approx 1.678$. Hence, as far as α is small, the leakiness term can improve the training of NN’s trying to learn the mode via the LGK loss, while introducing only an α -suppressed bias.

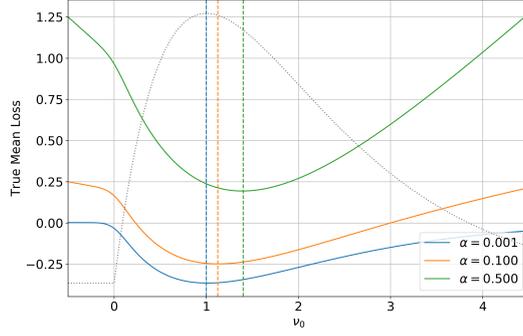


Figure 2. The true mean LGK losses with different α 's for the gamma distribution. The minimum of each mean loss function is shown with a dotted vertical line. The true mode of the distribution is at $y = 1$, and its median at $y \approx 1.678$.

3.4 Standard deviation

The three parameters mentioned so far are useful for characterizing the central tendency (i.e. some representative value) of the detector's response. However, also important for characterizing the detector are measures of the dispersion of the response distribution. Such quantities allow for parametrization of the precision of the detector's measurements.

For example, one may choose to characterize the dispersion of a detector's response using the standard deviation $\sigma(y|x)$ of its response. This can be parametrized using two NN's each trained using ℓ_{MSE} . In particular, a modified dataset can be combined with ℓ_{MSE} to learn any moment of the distribution of $p(y|x)$. This is accomplished by training with examples drawn from X and Y^n instead of the usual X and Y . One network trained with $n = 1$ (same as section 3.1) and another network trained with $n = 2$ can be combined to parametrize the standard deviation of the response:

$$\sigma(y|x) \equiv \sqrt{\mathbb{E}[y^2|x] - \mathbb{E}[y|x]^2}.$$

While a network trained to learn the second moment of y works well (similarly to learning the mean), the standard deviation estimated in this way is unstable because it requires taking the sum in quadrature of two estimates. Numerical instabilities in the required fine cancellation of large numbers results in a poor estimate of $\sigma(y|x)$.

3.5 Quantile

Another useful measure for the dispersion of the detector response is the quantile. As with the median, the advantage of statistics like the interquartile range over the standard deviation is the reduces sensitivity to outliers. A particular quantile of the response distribution can also be used to parametrize the upper/lower limits on the detector's output.

One can learn arbitrary quantiles of $p(y|x)$ by extending the absolute error ℓ_{AE} to the generalized absolute-error (GAE) [6, 7]:

$$\ell_{\text{GAE}}(y, \nu(x)) = \begin{cases} q|y - \nu(x)| & y \geq \nu(x) \\ (1 - q)|y - \nu(x)| & y < \nu(x) \end{cases} \quad (3.4)$$

where $q \in [0, 1]$ represents the quantile that one wants to learn. This can also be extended to be smooth [8, 9], simultaneously learn multiple quantiles [10], as well as other variations and extensions.

The mean training loss using GAE is minimized when:

$$\frac{\int_{v(x)}^{\infty} p(y|x) dy}{\int_{-\infty}^{v(x)} p(y|x) dy} = \frac{1-q}{q}$$

which precisely requires that $v(x)$ is the $100q\%$ quantile of the distribution $p(y|x)$.

4 Numerical examples

This section uses two examples from collider high-energy physics in order to illustrate the concepts presented in earlier sections.

Monte Carlo samples $\{(x_i, y_i)\}$ are generated following distributions $p(y|x)$ that are motivated by realistic detectors. In particular, the first example approximates the calorimeter response of a hadronic jet and the second example is motivated by the momentum reconstruction of muons. The jet energy is used to demonstrate the power of mode learning and the muon momentum reconstruction illustrates the potential of quantile learning. In both cases, the prior distribution $p(x)$ is chosen to be uniform and the goal is to compare the mode or quantiles to the analytic result based on $p(y|x)$.

4.1 Calorimeter and jet energy response

The first example is a calorimeter used to measure hadronic jets in collider experiments. Jets are collimated sprays of particles resulting from quarks and gluons produced at high energy. Details of the ATLAS and CMS jet calibrations can be found in refs. [11–15].

Because the calorimeter measures the energy of a large number of particles produced in a hadronic particle shower, the jet energy response is approximately Gaussian. However, jet energy measurements also typically demonstrate non-Gaussian tails due to various detector effects such as inactive material, punch-through beyond the calorimeter, etc. In addition, jet reconstruction algorithms often place a minimum transverse energy threshold to reduce the high rate of noise jets resulting from calorimeter noise, random combinations of low-energy particles, or jets from additional nearly simultaneous collisions (pileup). To illustrate the impact of these effects, the example considered in this section uses a truncated Crystal Ball distribution:

$$p(y|x) = \begin{cases} 0 & \text{if } y < E_{\min} \\ A_1 \left(\frac{m}{\beta} - \beta - \frac{y-\mu(x)}{\sigma(x)} \right)^{-m} & \text{if } E_{\min} < y < E_{\text{tail}} \\ A_2 \exp\left(-\frac{(y-\mu(x))^2}{2\sigma(x)^2}\right) & \text{if } y > E_{\text{tail}} \end{cases} \quad (4.1)$$

where x is the true jet energy, y is the reconstructed jet energy, E_{\min} is the minimum jet energy imposed by the reconstruction algorithm, and $A_{1,2}$ are normalization factors. E_{\min} is set to 10 in this study. β , m , and $E_{\text{tail}} \equiv \mu(x) - \beta\sigma(x)$ are parameters of the Crystal Ball distribution. β and m are set to 1.5 and 2, respectively. These values are chosen such that the threshold effects and the

non-Gaussian tails have a non-trivial impact on the mean energy response. The units are chosen to be typical for a calorimeter measurement in GeV.

The parameters $\mu(x)$ and $\sigma(x)$ characterize the central tendency and spread of the jet response. The reconstructed jet energy is typically smaller than the true jet energy because hadronic calorimeters are composed of layers of active and inactive (absorber) material and the latter do not collect any signal. For illustration, the functional form of μ is taken from ref. [16] and motivated by the ATLAS and CMS calorimeter jet energy response:

$$\mu(x) = x + 5 \ln\left(e^{x_0/10} + e^{x/10}\right) - 5 \ln\left(1 + e^{(x+x_0)/10}\right), \quad (4.2)$$

where x_0 is a turn-on parameter for the response function, set to 25 in this toy example. The jet energy resolution function $\sigma(x)$ uses the a typical parameterization taking into account a noise term (N), stochastic term (S), and constant term (C):

$$\sigma(x) = N \oplus S\sqrt{x} \oplus Cx, \quad (4.3)$$

where N represents energy-independent effects such as detector noise and pileup interactions, $S\sqrt{x}$ represents the Poisson-distributed, statistical fluctuations, and Cx represents fluctuations proportional to the true jet energy x , such as energy lost in the passive absorber volume. For the example presented here, $N = 4$, $S = 0.6$, and $C = 0.05$. These values roughly reproduce the transverse energy resolutions reported by ATLAS and CMS, with a relative resolution of approximately 25% (10%) at transverse energy of 20 GeV (100 GeV).

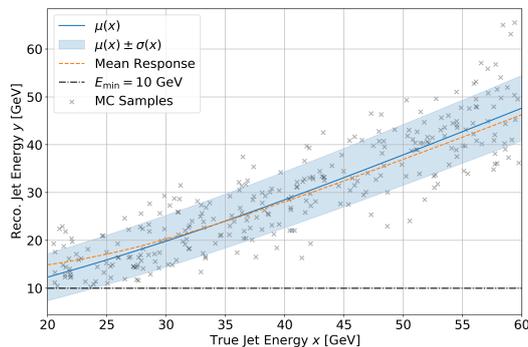


Figure 3. Jet energy response similar to those at ATLAS and CMS. 300 random MC samples are plotted together.

For training, 500,000 examples are generated using a uniform distribution in x and then the response distribution defined by eqs. (4.1), (4.2), and 4.3. These data are shown in figure 3 alongside the analytic values of μ and σ . Note that some low energy jets are lost because of the minimum energy cut at $E_{\min} = 10$ GeV, particularly below $x \approx 30$ GeV. As a result, the sample mean response is above the true value of $\mu(x)$. In the high-energy regime, threshold effect diminishes, and the non-Gaussian tail causes the mean response to be lower than $\mu(x)$. Therefore, if one were to characterize the jet energy response of the calorimeter, the mean response may not be an appropriate choice of parametrization.

To show the benefit of mode learning, two NN's are trained using the same dataset, one with the MSE loss (for the mean, section 3.1) and one with the LGK loss (for the mode, section 3.3).

The inputs to the networks are the true jet energy x and an additional feature x^3 , added to ease the learning of a non-linear function. The two neural networks have identical architectures: an input layer with two nodes, a single hidden layer with 50 nodes with parametric rectified linear unit (PReLU) [17] as activation functions, and an one-dimensional output layer. Both networks are implemented in Keras [18] using the Tensorflow backend [19] and trained using the ADAM optimization method [20], until their mean training losses converge. This process is repeated 25 times to obtain a set of 25 NN's parametrizing the mean and another set parametrizing the mode.

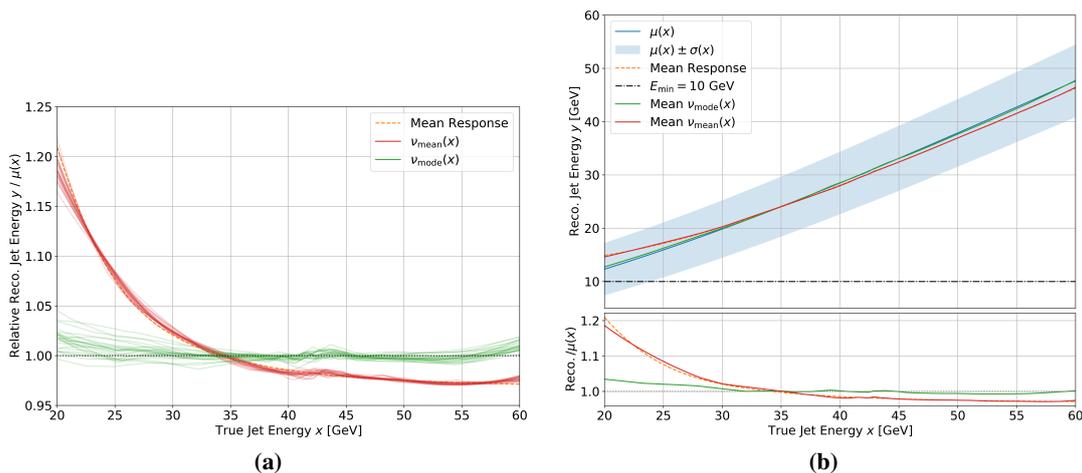


Figure 4. (a) All NN's trained to parametrize the jet energy response: 25 ν_{mean} 's and ν_{mode} 's each. ν_{mean} 's are trained with the MSE loss, and ν_{mode} 's are trained with the mean LGK loss with $h = 0.1$ and $\alpha = 10^{-4}$. (b) True parameters and NN-parametrizations of the jet response. The red and green lines show the mean of 25 ν_{mean} 's and ν_{mode} 's respectively.

Figure 4a shows the final results obtained from all the NN's. The ν_{mode} 's trained with the LGK loss generally follow the mode of the response $\mu(x)$, while the ν_{mean} 's trained with the MSE loss follows the mean response (distorted from $\mu(x)$). To reduce variance from initialization and training processes, the average of 25 NN's of each type are combined. These results are shown in figure 4b. The mode-learning NN's are clearly more appropriate, if one were to parametrize the jet energy response robust to the minimum-energy cutoff and the non-Gaussian tails.

4.2 Muon momentum resolution

This second example is used to illustrate the power of quantile learning to characterize the spread of the response distribution. For this purpose, muon momentum reconstruction provides the motivation. Characterizing the muon momentum resolution is important for resonance searches [21, 22] and has both Gaussian and non-Gaussian components. Further details of the muon momentum reconstruction can be found in refs. [23–26].

As in the previous example, a Crystal Ball distribution is used to model the resolution for the reconstructed muon momentum. In particular, $y|x = 1 + \sigma(x)z$ where z follows a standard Crystal Ball distribution (eq. (4.1) with $\mu = 0$, $\sigma = 1$, $E_{\text{min}} = -\infty$) with $m = 2$. The resolution function is

chosen to approximately match the ATLAS and CMS momentum resolution:

$$\sigma(x) = 0.03 \left(\frac{1}{4} e^{20 \log_{10}(\log_{10}(x)) - 4} + 1 \right), \quad (4.4)$$

where x is the momentum, given in units of GeV. The distribution of $\log_{10}(x)$ is sampled uniformly for training between $\log_{10}(7)$ and 2. This distribution has the property that it is well-described by a Gaussian when $|x - 1|$ is smaller than the Crystal Ball β parameter and otherwise poorly described by a Gaussian. Furthermore, the heavy tail makes the standard deviation diverge. The goal is to show that learning the 68% inter-quantile range is a robust estimate of the core Gaussian standard deviation and is insensitive to the heavy tails governed by β .

To learn the 68% inter-quantile range, two neural networks are trained to learn the 84% and 16% quantiles using the GAE loss (section 3.5). Half of the difference between these NN outputs should approximate the Gaussian core standard deviation. Figure 5a shows the learned value of the 68% inter-quantile range compared with $\sigma(x)$ as a function of x for $\beta = 2$. Even though there are heavy tails, the quantile regression is able to accurately model the core Gaussian standard deviation. Figure 5b demonstrates that the GAE loss produces an estimate that is robust to variations in β . Smaller values of β correspond to heavier tails. For comparison, the sample standard deviation is shown alongside the estimate of the NN to show the impact of the heavy tails on a naive estimate of the standard deviation. As $\beta \rightarrow \infty$, the heavy tails are suppressed and the sample standard deviation approaches the core Gaussian standard deviation.

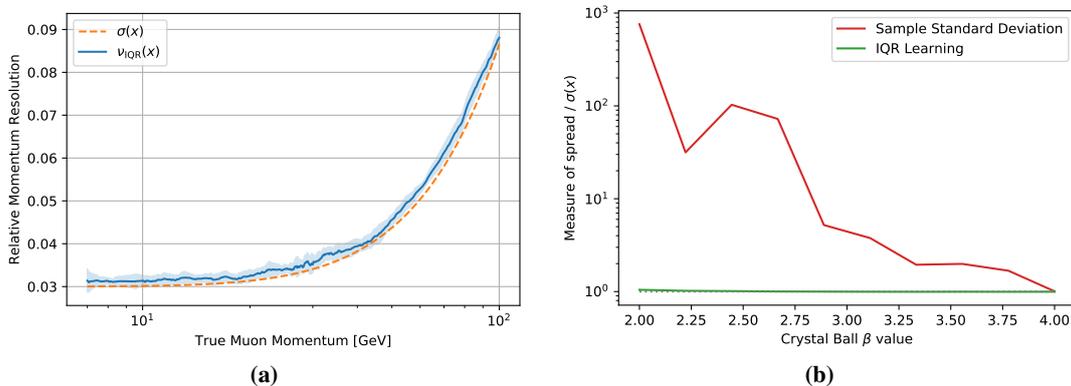


Figure 5. (a) The median (central value) and 68% inter-quantile range (error band) using the 25 NN's trained to parametrize the 68% inter-quantile range of the muon momentum response distribution with $\beta = 2$. The NN's are optimized using the GAE loss. (b) The measured spread divided by $\sigma(x)$ as a function of β . For the sample standard deviation, this is the sample standard deviation of $(y - 1)/\sigma(x)$ over all (x, y) values. The corresponding quantity for the NN is the average value of $\frac{1}{2}(v_{\text{GAE},q=0.16}(x) - v_{\text{GAE},q=0.84}(x))/\sigma(x)$.

5 Discussion

NN's can provide data-driven parameterizations of detector responses, free of *ad hoc* modeling assumptions. These methods can be naturally extended to include high-dimensional inputs and outputs. Such NN-based parameterizations of an experiment can serve many purposes. One

important example is the development of fast simulation tools like DELPHES 3 [27], where the detector response of all objects are parameterized. NN-based parameterizations may provide more accurate models and require less person-hours to derive. This is especially important for feasibility studies of future experiments where detailed studies are often not possible.

The various loss functions may also be useful more generally, in order for the neural network training to be more tailored to the particular task and underlying data distribution. In particular, proposals to use generative NN's for calorimeter simulations [28–35] are often trained with loss functions to include an absolute error term $\ell(E, E') = |E - E'|$ for the true and reconstructed (generated) energies [28–30]. While this does encourage the reconstructed energy generally towards the center (median) of the response distribution, using the LGK loss and regressing towards the peak of the distribution may be more appropriate.

Another important example where response parameterizations are necessary is detector calibrations. Supposing that x is the true value and y is the measured value of some observable, the most basic calibration approach would be to calibrate y to $\nu(y)$, where ν is a NN trained to learn the central tendency of x given y . One problem with this approach is that the calibrated value depends on the distribution of x used to train the NN [16]. An alternative method that is used by ATLAS and CMS for jet calibrations is known as *numerical inversion*, where the response parameterization of y given x is inverted: $y \mapsto \nu^{-1}(y)$. Traditionally, this was done with fits and interpolation functions, but it can be generalized to neural networks where it can naturally depend on many features (*generalized numerical inversion*) [36]. As the non-NN approaches typically use the mode, generalized numerical inversion with the LGK loss may be more appropriate than the MSE loss studied thus far. One may further extend generalized numerical inversion to also include confidence intervals by using the quantile learning loss GAE in addition to only the measures of central tendency.

6 Conclusions

This paper has reviewed various strategies for machine learning different summary statistics of a detector response distribution $p(y|x)$ with neural networks including the mean, median, and mode for central tendency and the standard deviation and quantiles for measures of spread. The summary statistic is linked to the loss function used in the neural network training. While most of the loss functions have been combined with neural networks before, the authors believe that mode learning neural networks are presented here for the first time. This is a particularly important case for detector parameterizations where the most likely value is often used instead of the mean or median to characterize the central tendency of the response distribution. All of the methods presented here will expand the toolkit of available methods for detector characterization with a wide variety of applications in experimental particle physics and beyond.

Note added. During the final preparation of this paper, a study from the CMS collaboration on b -jet calibration became public [37]. This note used the Huber loss to characterize the central tendency and quantile learning to parameterize the resolution. As far as the authors are aware, this is the first use of the latter technique in high energy physics.

Acknowledgments

This work was supported by the U.S. Department of Energy, Office of Science under contracts DE-AC02-05CH11231 and DE-AC02-76SF00515.

References

- [1] CMS collaboration, *Performance of Photon Reconstruction and Identification with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV*, 2015 *JINST* **10** P08010 [[arXiv:1502.02702](#)].
- [2] ATLAS collaboration, *Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data*, 2019 *JINST* **14** P03017 [[arXiv:1812.03848](#)].
- [3] P.J. Huber, *Robust estimation of a location parameter*, *Ann. Math. Statist.* **35** (1964) 73.
- [4] M. Lee, *Mode Regression*, *J. Econometrics* **42** (1986) 337.
- [5] W. Yao and L. Li, *A New Regression Model: Modal Linear Regression*, *Scand. J. Stat.* **41** (2013) 656.
- [6] R. Koenker and G. Bassett, *Regression Quantiles*, *Econometrica* **46** (1978) 33.
- [7] J.W. Taylor, *A Quantile Regression Neural Network Approach To Estimating the Conditional Density of Multiperiod Returns*, *J. Forecasting* **19** (2000) 299.
- [8] S. Zheng, *Gradient descent algorithms for quantile regression with smooth approximation*, *Int. J. Mach. Learn. Cybern.* **2** (2011) 191.
- [9] K. Hatalis, A.J. Lamadrid, K. Scheinberg and S. Kishore, *Smooth Pinball Neural Network for Probabilistic Forecasting of Wind Power*, [arXiv:1710.01720](#).
- [10] N. Tagasovska and D. Lopez-Paz, *Frequentist uncertainty estimates for deep learning*, [arXiv:1811.00908](#).
- [11] ATLAS collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **73** (2013) 2304 [[arXiv:1112.6426](#)].
- [12] ATLAS collaboration, *Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **75** (2015) 17 [[arXiv:1406.0076](#)].
- [13] ATLAS collaboration, *Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **96** (2017) 072002 [[arXiv:1703.09665](#)].
- [14] CMS collaboration, *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, 2017 *JINST* **12** P02014 [[arXiv:1607.03663](#)].
- [15] CMS collaboration, *Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS*, 2011 *JINST* **6** P11002 [[arXiv:1107.4277](#)].
- [16] A. Cukierman and B. Nachman, *Mathematical Properties of Numerical Inversion for Jet Calibrations*, *Nucl. Instrum. Meth. A* **858** (2017) 1 [[arXiv:1609.05195](#)].
- [17] K. He, X. Zhang, S. Ren and J. Sun, *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*, [arXiv:1502.01852](#).
- [18] F. Chollet, *Keras*, (2017) <https://github.com/fchollet/keras>.

- [19] M. Abadi et al., *Tensorflow: A system for large-scale machine learning*, in proceedings of the *12th USENIX Symposium on Operating Systems Design and Implementation*, Savannah, GA, U.S.A., 2–4 November 2016, volume 16, USENIX Association, Berkeley CA U.S.A. (2016), pp. 265–283.
- [20] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, in proceedings of the *3rd International Conference on Learning Representations (ICLR 2015)*, San Diego, CA, U.S.A., 7–9 May 2015, [arXiv:1412.6980](#).
- [21] ATLAS collaboration, *Search for high-mass dilepton resonances using 139 fb⁻¹ of pp collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Lett. B* **796** (2019) 68 [[arXiv:1903.06248](#)].
- [22] CMS collaboration, *Search for high-mass resonances in dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **06** (2018) 120 [[arXiv:1803.06292](#)].
- [23] ATLAS collaboration, *Muon reconstruction performance of the ATLAS detector in proton-proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **76** (2016) 292 [[arXiv:1603.05598](#)].
- [24] ATLAS collaboration, *Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data*, *Eur. Phys. J. C* **74** (2014) 3130 [[arXiv:1407.3935](#)].
- [25] CMS collaboration, *Performance of CMS Muon Reconstruction in pp Collision Events at $\sqrt{s} = 7$ TeV, 2012 JINST 7 P10002* [[arXiv:1206.4071](#)].
- [26] CMS collaboration, *Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV*, *2018 JINST 13 P06015* [[arXiv:1804.04528](#)].
- [27] DELPHES 3 collaboration, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)].
- [28] M. Paganini, L. de Oliveira and B. Nachman, *Accelerating Science with Generative Adversarial Networks: An Application to 3D Particle Showers in Multilayer Calorimeters*, *Phys. Rev. Lett.* **120** (2018) 042003 [[arXiv:1705.02355](#)].
- [29] M. Paganini, L. de Oliveira and B. Nachman, *CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks*, *Phys. Rev. D* **97** (2018) 014021 [[arXiv:1712.10321](#)].
- [30] ATLAS collaboration, *Deep generative models for fast shower simulation in ATLAS*, *ATL-SOFT-PUB-2018-001* (2018).
- [31] F. Carminati, A. Gheata, G. Khattak, P. Mendez Lorenzo, S. Sharan and S. Vallecorsa, *Three dimensional Generative Adversarial Networks for fast simulation*, *J. Phys. Conf. Ser.* **1085** (2018) 032016.
- [32] V. Chekalina, E. Orlova, F. Ratnikov, D. Ulyanov, A. Ustyuzhanin and E. Zakharov, *Generative Models for Fast Calorimeter Simulation: the LHCb case*, in proceedings of the *23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2018)*, Sofia, Bulgaria, 9–13 July 2018, *EPJ Web Conf.* **214** (2019) 02034 [[arXiv:1812.01319](#)].
- [33] M. Erdmann, J. Glombitza and T. Quast, *Precise simulation of electromagnetic calorimeter showers using a Wasserstein Generative Adversarial Network*, *Comput. Softw. Big Sci.* **3** (2019) 4 [[arXiv:1807.01954](#)].
- [34] L. de Oliveira, M. Paganini and B. Nachman, *Controlling Physical Attributes in GAN-Accelerated Simulation of Electromagnetic Calorimeters*, *J. Phys. Conf. Ser.* **1085** (2018) 042017 [[arXiv:1711.08813](#)].

- [35] M. Erdmann, L. Geiger, J. Glombitza and D. Schmidt, *Generating and refining particle detector simulations using the Wasserstein distance in adversarial networks*, *Comput. Softw. Big Sci.* **2** (2018) 4 [[arXiv:1802.03325](#)].
- [36] ATLAS collaboration, *Generalized Numerical Inversion: A Neural Network Approach to Jet Calibration*, [ATL-PHYS-PUB-2018-013](#) (2018).
- [37] CMS collaboration, *A deep neural network for simultaneous estimation of b quark energy and resolution*, [CMS-PAS-HIG-18-027](#) (2019).