
Loss function to optimise signal significance in particle physics

Jai Bardhan¹
jai.bardhan90@gmail.com

Cyrin Neeraj¹
cyrin.neeraj@research.iiit.ac.in

Subhadip Mitra¹
subhadip.mitra@iiit.ac.in

Tanumoy Mandal²
tanumoy.mandal@iisertvm.ac.in

¹Center for Computational Natural Science and Bioinformatics,
International Institute of Information Technology, Hyderabad

²Indian Institute of Science Education and Research, Thiruvananthapuram

Abstract

We construct a surrogate loss to directly optimise the significance metric used in particle physics. We evaluate our loss function for a simple event classification task using a linear model and show that it produces decision boundaries that change according to the cross sections of the processes involved. We find that the models trained with the new loss have higher signal efficiency for similar values of estimated signal significance compared to ones trained with a cross-entropy loss, showing promise to improve sensitivity of particle physics searches at colliders.

1 Introduction

Particle physics experiments at the Large Hadron Collider (LHC) rely heavily on multivariate classifiers to isolate signals from backgrounds. These investigations are generally of two types: 1) measuring known processes/properties with improving precision and checking for anomalies [i.e., departures from the predictions of the Standard Model] and 2) looking for new processes (like looking for hypothetical particles). In most cases, the need for multivariate classifiers comes from the sporadic nature of the signal compared to the background. Generally, the signal plus background hypothesis (H_1) is tested against the null or background-only hypothesis (H_0), and the disagreement between them is expressed in terms of a p value. An equivalent interpretation of the p value is the significance score (Z) defined such that a Gaussian-distributed variable found Z standard deviations away from its mean has a tail-distribution probability equal to p [2]. Most sensitivity studies commonly use a simple approximation of the median Z score as a measure of the estimated signal significance,

$$Z \approx N_s / \sqrt{N_b} \tag{1}$$

where N_s and N_b are the estimated numbers of signal and background events, respectively. (In the rest of the paper, we shall refer to the $\text{med}[Z]$ score as just the Z score.)

In this paper, we attempt to construct a loss function whose minimisation can directly enhance the experimental sensitivity. Our motivation comes from two observations. First, not all event rates are equal; some scattering processes have higher probabilities [parameterised as cross sections (σ), calculable from theory. The number of events from a process ($N_{s,b}$) is calculated as $N_{s,b} = \sigma_{s,b} \mathcal{L}$, where \mathcal{L} is the experimental luminosity] than others. Since a basic binary cross-entropy (BCE) loss

treats all events equally, it might wrongly classify events of some types which is more detrimental to the classifier performance than the other. Second, a loss that optimises the signal-to-background ratio ($r = N_s/N_b$) may not necessarily maximise the significance as the Z score depends on the ratio and the absolute number (i.e., set size) of signal or background events ($Z \approx \sqrt{N_s} \cdot \sqrt{r} = r\sqrt{N_b}$). Hence we ask, can we derive a loss function that maximises the Z score directly? The Z score described in Eq. (1) is a set function. Therefore, we define a surrogate loss function using the Lovász extension [1] to maximise it directly. We evaluate this loss with pseudo-data mimicking a typical event classification task using a linear model and compare the decision boundaries to that model trained on a BCE loss. We also compare the performance of models trained on a BCE loss.

2 Constructing the loss: submodularity and Lovász’s extension

We must consider some points before constructing a loss function based on the Z score. First, since the Z score is not a differentiable function (it depends upon discrete quantities), it needs a smooth interpolation. Second, the metric operates on datasets instead of individual samples—particularly, count data. Therefore, we must either develop a method to directly optimise the set function or assign contributions to specific samples within the set to optimise.

We look for a smooth submodular function. A submodular function is a function that captures the concept of diminishing returns. It is defined on sets and has a property similar to concavity. Formally, submodularity can be defined as:

Submodularity: A set function $\Delta : \{0, 1\}^p \rightarrow \mathbb{R}$ is submodular if for all sets $A, B \in \{0, 1\}^p$,

$$\Delta(A) + \Delta(B) \geq \Delta(A \cup B) + \Delta(A \cap B), \quad (2)$$

or, equivalently for $B \subseteq A$ and $i \notin A, i \notin B$,

$$\Delta(A \cup \{i\}) - \Delta(A) \leq \Delta(B \cup \{i\}) - \Delta(B). \quad (3)$$

The submodular functions can be optimised using greedy optimisation techniques, and it is to find optimal solutions in polynomial times. However, these discrete optimisation techniques cannot be used directly without a gradient. The Lovász extension allows us to associate a continuous, convex function with any submodular function:

Lovász extension: For a set function $\Delta : \{0, 1\}^p \rightarrow \mathbb{R}$, the Lovász extension $\bar{\Delta} : [0, 1]^p \rightarrow \mathbb{R}$ is defined as

$$\bar{\Delta} : \mathbf{m} \in \mathbb{R}^p \mapsto \sum_{i=1}^p m_i g_i(\mathbf{m}) \quad (4)$$

where $\mathbf{m} \in \mathbb{R}_+^p$ is the vector of errors (which we discuss in the next section), $g_i(\mathbf{m}) = \Delta(\{\pi_1, \dots, \pi_i\}) - \Delta(\{\pi_i, \dots, \pi_{i-1}\})$ and π is a permutation ordering the components of \mathbf{m} in decreasing order, i.e., $x_{\pi_1} \geq x_{\pi_2} \geq \dots \geq x_{\pi_p}$ [1]. For the Lovász extension to be applicable, the set function must be submodular.

Additionally, the Lovász extension of a submodular function preserves submodularity, i.e., the extension evaluated at the points of the hypercube still follows submodularity. Using the Lovász extension, we can directly compute the tight convex closure of a submodular function within polynomial time [$\mathcal{O}(p \log(p))$ time complexity]. This convex extension is amenable to a host of efficient optimisation methods, especially gradient-based approaches.

2.1 Submodularity of the Z score

We modify the Z score defined in Eq. (1) by setting $N_b = \epsilon + N_b$ to prevent the term from diverging at $N_b = 0$. We also add a minus sign to make it suitable for minimisation. The surrogate set function for the Z -score is then given by,

$$\Delta_Z(y, \tilde{y}, \boldsymbol{\sigma}, \mathbf{v}, \mathcal{L}) = \sum_{i \in S} \frac{\sigma_i \mathcal{L}}{\sqrt{\epsilon}} - \frac{\sum_{i \in S} \frac{v_i - n_i}{v_i} \sigma_i \mathcal{L}}{\sqrt{\epsilon + \sum_{i \in B} \frac{p_i}{v_i} \sigma_i \mathcal{L}}} \quad (5)$$

where y stands for the ground-truth labels of a set of events and \tilde{y} for the labels predicted by a model on the set; P_y and $P_{\tilde{y}}$ represent the set of positive labels and positive predictions and v_i is the number of events of process type i where $i \in S \cup B$, where S, B are the set of signal and background processes. The constant term $(\sum_{i \in S} \sigma_i \mathcal{L} / \sqrt{\epsilon})$ is added to ensure $\Delta_Z(\emptyset) = \mathbf{0}$. The function Δ_Z is submodular on the set of mispredictions (n, p) , where n is the number of false negatives, and p is the number of false positives (which can be calculated from P_y and $P_{\tilde{y}}$). The proof is presented in Appendix A. Even though here we consider only one signal process, the proof can be trivially generalised to the cases with multiple signal processes (n_1, n_2, \dots, n_r) as well. From Eq. (4), we see that for Δ_Z to be a loss function, the vector \mathbf{m} must be the error vector in the prediction; Δ_Z , then naturally emerges as the surrogate loss to optimise the Z score.

Choice of \mathbf{m} error: There are a few choices in the literature for modelling the error vector m . To illustrate the working of the Z -score loss for this paper, we pick a hinge (Max Margin) error similar to Ref. [4]. The labels are considered signed ($y_i \in \{-1, 1\}$). The model outputs a score $F_i(x)$ for each sample x . The error is given by the hinge loss,

$$m_i = \max(1 - F_i(x)y_i, 0), \quad y_i \in \{-1, 1\}. \quad (6)$$

The vector \mathbf{m} could also be modelled as a sigmoid error or cross-entropy error, for example. We plot the Z -score loss landscape for all these errors in the appendix for the toy problem (described below) in Appendix B.

There is only one free parameter in our loss: ϵ . Other quantities like σ_i and \mathcal{L} are set by the process under consideration (i.e., the particular classification task) and the collider experiment. Assuming we perform the classification for rare signals, we set $\epsilon = \sigma_s \mathcal{L}$, the theoretically predicted number of signal events (which is also the maximum number of estimated signal events) for testing the loss. Algorithm 1 provides a simple pseudocode to calculate the gradient $g(\mathbf{m})$ from Eq. (4) using Eq. (5) as the loss.

3 Testing the loss

We analyse the loss function with a simple toy problem which can be easily mapped to the problem of event classification at the LHC. Our goal is to separate the signal (s) from background events using a linear classifier in the presence of multiple (say, two, $b1$ and $b2$) dominant background processes, as is usually the case. The datasets are modelled as normal distributions in two features, x_1 and x_2 which can be thought of as the kinematic features of the actual events. We generated roughly 50000 points for each process and the optimisation was done in batches using RAdam optimiser. We train the linear classifier using the BCE loss and $\bar{\Delta}_Z$ with the hinge error for the following two test cases:

Case 1: $\sigma_{b1} = 1 \text{ fb}, \sigma_{b2} = 100 \text{ fb}; \sigma_s = 0.1 \text{ fb}$.

Case 2: $\sigma_{b1} = 100 \text{ fb}, \sigma_{b2} = 1 \text{ fb}; \sigma_s = 0.1 \text{ fb}$.

with $\mathcal{L} = 3000 \text{ fb}^{-1}$. We show the data distributions in Fig. 1. Since $\bar{\Delta}_Z$ has the event rate (the true probabilities) information, we expect the decision boundaries to be different for the two test cases—eliminating more events from the larger background will give better significance scores, which $\bar{\Delta}_Z$ is designed to optimise. Fig. 1 confirms this.

Performance: To compare the performance of the BCE- and $\bar{\Delta}_Z$ -trained models, we show the results of some scans in Case 1 in Fig. 2: the estimated Z score value for different model thresholds (u , varying which essentially translates the decision boundaries in Fig. 1 along the axes) and the variation of the estimated Z score with the signal efficiency ($\epsilon(u)$, the fraction of signal events retained for the threshold u) against the Z score obtained. For the scan, we demand $\epsilon(u) \geq 0.05$ and

Algorithm 1 Gradient of Lovász Z loss $\bar{\Delta}_Z$

Require: vector of errors $\mathbf{m} \in \mathbb{R}_+^p$, ground truth labels δ , sample weights $w = \{w_1, w_2, \dots, w_p\}$ calculated from σ_i and counts.

Ensure: $g(\mathbf{m})$ gradient of $\bar{\Delta}_Z$ from Equation (4)

- 1: $\pi \leftarrow$ decreasing sort permutation for \mathbf{m}
- 2: $\delta_\pi \leftarrow (\delta_{\pi_i})_{i \in [1, p]}$
- 3: **numerator** $\leftarrow 1 - \text{cumulative_sum}(\delta_\pi) w$
- 4: **denominator** $\leftarrow 1 + \text{cumulative_sum}(1 - \delta_\pi) w$
- 5: $g \leftarrow \sigma - \text{numerator} / \sqrt{\text{denominator}}$
- 6: **if** $p > 1$ **then**
- 7: $g[2 : p] \leftarrow g[2 : p] - g[1 : p - 1]$
- 8: **end if**
- 9: **return** g_π

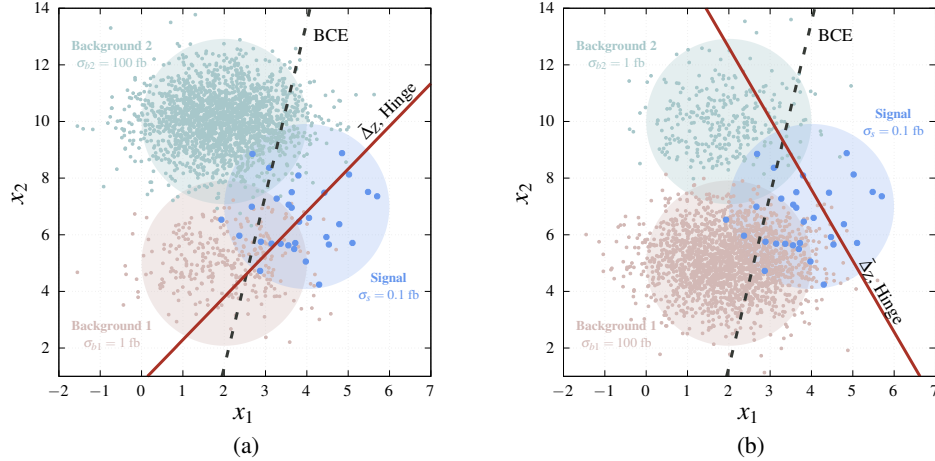


Figure 1: Decision boundaries of the linear classifier when trained with the $\bar{\Delta}_Z$ loss with hinge error for (a) Case 1 and (b) Case 2 (see Section 3). When trained with $\bar{\Delta}_Z$, the classifier prioritises reducing the background with the larger cross section.




Figure 2: (a) The estimated Z score for the entire range of the linear model threshold, u . (b) The distribution of the Z score with signal efficiency, the fraction of signal events retained. Both quantities are functions of u . The model trained with $\bar{\Delta}_Z$ reaches the maximum Z score for higher values of signal efficiency than that with the BCE loss. (c) Class efficiencies vs. u when trained with the $\bar{\Delta}_Z$ loss with the hinge error. With increasing u , the larger background is eliminated first. For very high u the drop in the subdominant background is less steeper than the signal, leading to the drop in $Z(u)$ in (a) for higher thresholds.

the number of background events beyond the threshold to be at least 5. Similar plots are obtained for Case 2 also. From the figure, we see that $\bar{\Delta}_z$ maximises the Z score for a higher signal efficiency than the BCE, i.e., where the estimated Z score peaks, the model retains more signal events than the BCE-trained model. (For the $\bar{\Delta}_z$ model, the estimated Z score drops for high values of u because there, for the datasets we consider, the major background is almost eliminated and further shifting the decision boundary reduces the minor background slower than the signal).

4 Conclusions and Outlook

In this article, we showed how a loss function for directly optimising the signal significance can be constructed. We obtained a surrogate for the median Z score, proved that it is a submodular set function and derived a loss function that can be used to train a multivariate model in batches using the Lovász extension. We showed that models trained with such a loss can cut the heavy background(s) more than the ones trained on the BCE loss while retaining more signal events (and thus showing the promise of enhancing experimental sensitivities).

Limitations and scope: While our results are promising, further tests are needed to fully characterise and understand the benefits and limitations of $\bar{\Delta}_Z$. Here, our choice of using a linear classifier on simple datasets was motivated by its simplicity and interpretability. However, for realistic characterisations, one has to look beyond the linear classifier (e.g., use a deep neural network) and consider a range of benchmark (new-physics) scenarios with different kinematics (features). For example, there could be multiple (more than two) major backgrounds with highly overlapping features or the signal size could be much smaller than the backgrounds (more than what we considered, as is the case in some heavy particle searches).

Finally, we note that while it is possible to introduce rate-dependent weights directly in the BCE loss, tuning them is an empirical task. The weights that yield the best performance need not be simply the rates of the processes. In contrast, $\bar{\Delta}_Z$ presents a natural way to include the rates (cross sections) as it is derived from the significance score used in collider searches.

5 Code availability

A PYTHON implementation for $\bar{\Delta}_Z$ is available at <https://github.com/Jai2500/Z-Score-Loss>.

6 Acknowledgments

J. B. would like to thank Arpan Dasgupta for valuable discussions on submodular functions.

References

- [1] Maxim Berman and Matthew B. Blaschko. Optimization of the jaccard index for image segmentation with the lovász hinge. *CoRR*, abs/1705.08790, 2017.
- [2] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C*, 71:1554, 2011. [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)].
- [3] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [4] Jiaqian Yu and Matthew Blaschko. The lovász hinge: A novel convex surrogate for submodular losses, 2015.

A Proof of sub-modularity of the Z score

For the proof, we take the scenario with a single signal process ($|S| = 1$, $\mathbf{n} = n_1 = n$) and a single background process ($|P| = 1$, $\mathbf{p} = p_1 = p$) to simplify the expressions. But the result can be easily extended to incorporate multiple signal and background processes due to the linearity of additional signal processes and background processes. We will also drop the luminosity term as that will not affect the core derivation.

For the proof, let us assume that we have two sets of misclassifications $C (n_C, p_C)$ and $D (n_D, p_D)$, such that $D \subseteq C$, i.e.,

$$D \subseteq C, \quad n_D \leq n_C, \quad p_D \leq p_C \quad (7)$$

The total number of events remains the same between C and D , i.e., v_S for signal and v_B for background, and only the misclassifications on the total set change.

To establish the proof, we need to show that the diminishing return property of submodularity holds under the addition of a new element $i \notin C$.

Case I: Adding false negatives $i \notin C$

We prove that Δ_Z is submodular under the addition of false negatives:

$$\Delta_Z(C \cup \{i\}) = \Delta_Z(n_C + 1, p_C) \quad (8)$$

$$= \frac{\sigma_S}{\sqrt{\epsilon}} - \frac{\frac{v_S - n_C - 1}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B}} \quad (9)$$

$$= \left(\frac{\sigma_S}{\sqrt{\epsilon}} - \frac{\frac{v_S - n_C}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B}} \right) + \frac{\frac{1}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B}} \quad (10)$$

$$= \Delta_Z(C) + \frac{\frac{1}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B}} \quad (11)$$

$$\Delta_Z(C \cup \{i\}) - \Delta_Z(C) = \frac{\frac{1}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B}} \quad (12)$$

Now since $D \subseteq C$, i.e., $p_D \leq p_C$, we see from Eq. (12),

$$\Delta_Z(C \cup \{i\}) - \Delta_Z(C) \leq \Delta_Z(D \cup \{i\}) - \Delta_Z(D), \quad i \text{ is a false negative} \quad (13)$$

Case II: Adding a false positive $i \notin C$

We prove that Δ_Z is submodular under the addition of false positives:

$$\Delta_Z(C \cup \{i\}) = \Delta_Z(n_C, p_C + 1) \quad (14)$$

$$= \frac{\sigma_S}{\sqrt{\epsilon}} - \frac{\frac{v_S - n_C}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B + \frac{1}{v_B} \sigma_B}} \quad (15)$$

$$(16)$$

Now we have,

$$\Delta_Z(C \cup \{i\}) - \Delta_Z(C) = \frac{\frac{v_S - n_C}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B}} - \frac{\frac{v_S - n_C}{v_S} \sigma_S}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B + \frac{1}{v_B} \sigma_B}} \quad (17)$$

$$= \underbrace{\left(\frac{v_S - n_C}{v_S} \sigma_S \right)}_{T_1} \underbrace{\left[\frac{1}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B}} - \frac{1}{\sqrt{\epsilon + \frac{p_C}{v_B} \sigma_B + \frac{1}{v_B} \sigma_B}} \right]}_{T_2} \quad (18)$$

Therefore it decomposes into a product of two terms. If we show that independently both of these terms are independently smaller for C than for D , we will have our result.

First consider T_1 , we have

$$n_C \geq n_D \quad (19)$$

$$= -n_C \leq -n_D \quad (20)$$

$$= \frac{v_S - n_C}{v_S} \leq \frac{v_S - n_D}{v_S} \quad (21)$$

Therefore, T_1 is indeed larger for D compared to C .

In order to check for term T_2 , we first simplify the expression and write $H_C = \epsilon + \frac{p_C}{v_B} \sigma_B$, ($H_C \geq H_D$). Now we can write term two as:

$$\frac{1}{\sqrt{H_C}} - \frac{1}{\sqrt{H_C + \frac{1}{v_B} \sigma_B}} \quad (22)$$

We move to a continuous relaxation of the term such that:

$$f(x) = \frac{1}{\sqrt{x}} - \frac{1}{\sqrt{x + \frac{1}{v_B} \sigma_B}} \quad (23)$$

$$f(H_C) = \frac{1}{\sqrt{H_C}} - \frac{1}{\sqrt{H_C + \frac{1}{v_B} \sigma_B}} \quad (24)$$

which is the same as Eq. (22). Now differentiating Eq. (23) with respect to x , we get:

$$\frac{d}{dx} f = \frac{1}{2} \left(\frac{1}{\left(x + \frac{1}{v_B} \sigma_B\right)^{\frac{3}{2}}} - \frac{1}{x^{\frac{3}{2}}} \right) \quad (25)$$

which will always be less than zero for $x > 0$. Thus since $\frac{d}{dx} f < 0$, we have that T_2 will be greater for D compared to C .

Now since both T_1 and T_2 is greater for D compared to C , we have

$$\Delta_Z(C \cup \{i\}) - \Delta_Z(C) \leq \Delta_Z(D \cup \{i\}) - \Delta_Z(D), \quad i \text{ is a false positive} \quad (26)$$

Therefore from **Case I** and **Case II**, we have shown that Δ_Z is submodular for all the possible cases and therefore is submodular for the set of misclassifications (\mathbf{n}, \mathbf{p}) .

B Error Functions

We require a loss function to handle any vector of errors $\mathbf{m} \in \mathbb{R}_+^p$ since we are working with continuous predictions, not only to discrete vectors of misclassifications in $\{0, 1\}^p$. We consider four cases for defining the vector of errors \mathbf{m} to construct the surrogate losses using the Lovasz extension.


1. **Hinge (Max Margin) Loss:** Following Ref. [4], we implement a hinge loss to compute the error in the prediction. The labels are considered signed ($y_i \in \{-1, 1\}$). The model outputs a score $F_i(x)$ for each sample x . The error is given by the hinge loss,

$$m_i = \max(1 - F_i(x)y_i, 0), \quad y_i \in \{-1, 1\}. \quad (27)$$


2. **Sigmoid Error:** Similar to Ref. [1], we also consider the sigmoid error. The model outputs a probability $F_i(x)$ for the sample x to be in the signal class. The error is given by

$$m_i = \begin{cases} 1 - F_i(x), & \text{if } y_i = 1, \\ F_i(x), & \text{otherwise.} \end{cases} \quad (28)$$


3. **Cross Entropy Error:** We also experiment with the BCE loss to measure the error m_i . This is similar to taking the logarithm of the error calculated in the Sigmoid Error. One could also interpret this as a form of weighted cross entropy where the weights are calculated based on the specific composition of the batch of events and misclassifications on that batch.




(a) Hinge Error



(b) Sigmoid Error



(c) CE Error



(d) Focal Error

Figure 3: Loss landscapes for the four error measures \mathbf{m} in Sec. § 2.1. The Z score loss is plotted with ground truth, $GT = [1, 0]$, $\sigma = [1, 10]$. The x, y axes denote the classifier output $(F_1(x), F_2(x))$. For the Hinge Error, the GT labels are converted to their signed equivalent.

4. **Focal Loss Error:** We also consider Focal Loss [3] as the measure for the error m_i . This loss function drives the network to focus on hard misclassified events.

Our formulation results in a convex loss with a global minimum, which we evaluate for a simple case for all the error metrics as visualized in Fig. 3.

C ROC Curves for Case 1.

We plot the ROC curves for experiments for Case 1 in Fig. 4. Case 2 gives similar results. Let N_{B1}, N_{B2} represent the total number of BG1 and BG2 events generated in the dataset. Let n_{B1}, n_{B2} represent the number of BG1 and BG2 events remaining after the threshold respectively. Let σ_{B1}, σ_{B2} represent the cross sections of process BG1 and BG2 respectively. The total background efficiency is given by,

$$\frac{n_{B1} + n_{B2}}{N_{B1} + N_{B2}},$$

and the true background efficiency is given by,

$$\frac{\left(\frac{n_{B1}}{N_{B1}}\right)\sigma_{B1} + \left(\frac{n_{B2}}{N_{B2}}\right)\sigma_{B2}}{\sigma_{B1} + \sigma_{B2}}.$$

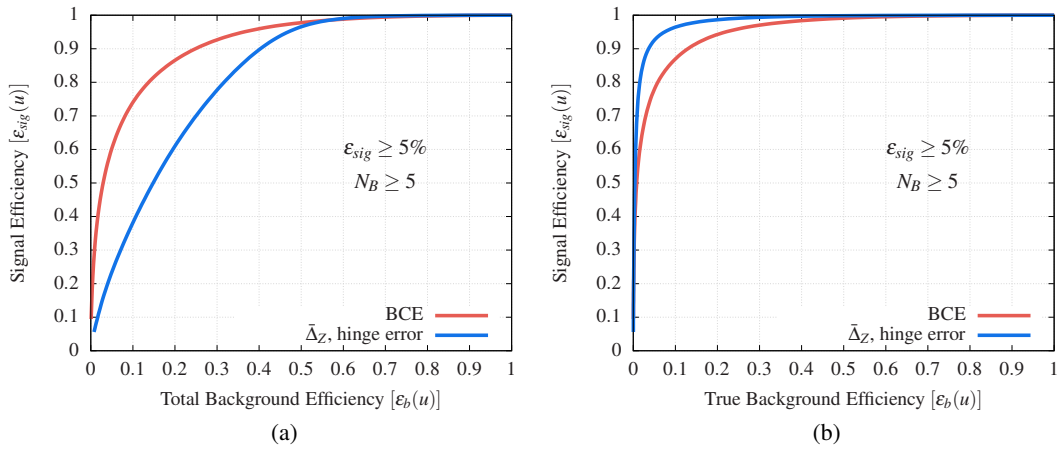


Figure 4: (a) ROC Curve for dataset (total) background efficiency vs signal efficiency. (b) ROC Curve for true background efficiency vs signal efficiency. The true background efficiency differs from the total background efficiency in that it accounts for the cross sections of the background processes. We observe that our loss performs better at removing background at a higher signal efficiency.