

PAPER • OPEN ACCESS

Machine Learning on sWeighted data

To cite this article: M Borisyak and N Kazeev 2020 *J. Phys.: Conf. Ser.* **1525** 012088

View the [article online](#) for updates and enhancements.

Recent citations

- [Muon identification for LHCb Run 3](#)
L. Anderlini *et al*



ECS **240th ECS Meeting**
Digital Meeting, Oct 10-14, 2021
We are going fully digital!
Attendees register for free!
REGISTER NOW

Machine Learning on sWeighted data

M Borisyak¹ N Kazeev^{1,2}

¹ Laboratory of Methods for Big Data Analysis, National Research University Higher School of Economics, 3 Kochnovsky Proezd, Moscow 125319, Russia

² Sapienza University of Rome, Piazzale Aldo Moro, 5, 00185 Roma RM, Italy

E-mail: nikita.kazeev@cern.ch

Abstract. Data analysis in high energy physics has to deal with data samples produced from different sources. One of the most widely used ways to unfold their contributions is the sPlot technique. It uses the results of a maximum likelihood fit to assign weights to events. Some weights produced by sPlot are by design negative. Negative weights make it difficult to apply machine learning methods. The loss function becomes unbounded. This leads to divergent neural network training. In this paper we propose a mathematically rigorous way to transform the weights obtained by sPlot into class probabilities conditioned on observables, thus enabling to apply any machine learning algorithm out-of-the-box.

1. Introduction

Data obtained by high energy physics experiments is often a mixture of signal and background events. Application of conventional classification methods to such data is a challenging task as obtaining precise label information is usually impossible. Instead, one can use a probabilistic approach by employing so-called discriminative variables (usually, the invariant mass), whose distributions for signal and background events are known in advance, or can be estimated by a maximum-likelihood fit. The sPlot technique [1] allows us to recover the distribution of target variables for signal events, given that target variables are statistically independent of discriminative ones. sPlot achieves that by assigning weights (sWeights) to all events, e.g. histogram of the target variable is obtained as the sum of these weights within each bin:

$$P(x \in Q | S) \approx \frac{1}{\sum_i w_i} \sum_i w_i \mathbb{I}[x_i \in Q], \quad (1)$$

where x_i are samples of the target variables, w_i are weights assigned by sPlot, Q denotes an arbitrary set and S a signal event. Figure 1 provides an illustration of sPlot technique. Derivation and detailed analysis of the sPlot are given in [1].

A notable property of the sPlot technique is that these weights can be negative. Typically, This does not create problems for low-dimensional analysis with simple models. Nonetheless, one can encounter difficulties fitting large-capacity models, e.g. deep neural networks. The prior work [2] describes such difficulties for the problem of signal/background separation and introduces two approaches for dealing with them. In this work, we expand the previous results to the problem of classification between two signal sources, samples of each are contaminated by different sources of background events.






Figure 1: Demonstration of the sPlot technique. The known distributions of a discriminative variable m (left) is used to weight the distribution of a target variable x for mixture of signal and background events (right). Adapted from [2].

The rest of the paper is structured as follows. In the next section, we briefly outline prior results from [2]. Section 3 introduces methods for signal-signal classification. Experimental results are presented in section 4.

2. Background

Consider the problem of training a classifier $f(x)$ to separate background (denoted as B) from signal (denoted as S) events given unlabeled sample $\{(x_i, m_i)\}_{i=1}^N$, where x denote target variables, m — discriminative variables. We assume that x and m are statistically independent:

$$p(x, m | S) = p(x | S)p(m | S); \quad (2)$$

$$p(x, m | B) = p(x | B)p(m | B); \quad (3)$$

and $p(m | S)$ and $p(m | B)$ are known in advance.

One can use a cross-entropy loss to train a classifier:

$$\mathcal{L}(f) = \mathbb{E}_{x, m \sim p(x, m)} w(m) \log f(x) + (1 - w(m)) \log(1 - f(x)) = \mathbb{E}_{x \sim p(x)} P(S | x) \log f(x) + P(B | x) \log(1 - f(x)), \quad (4)$$

where $p(x, m)$, $p(x)$ denote probability density functions and $w(m)$ the weights assigned by sPlot. The last equality is due to the property of sPlot:

$$\mathbb{E}_{x, m \sim p(x, m)} w(m) = P(S | x). \quad (5)$$

Proof of this statement can be found in [2].

Loss (4) can be approximated with a finite sample as follows:

$$L_1(f) = \sum_i w_i \log f(x_i) + (1 - w_i) \log(1 - f(x_i)). \quad (6)$$

This approach is quite popular in the high energy physics community [2]. One might notice, however, that some of the summands in (6) might not have a lower bound in case of negative w_i .

Therefore, if classifier $f(x)$ can isolate a point x_i with a negative corresponding weight w_i from the rest of the sample, loss (6) can be made arbitrarily low, which leads to quick overfitting. This is often possible for complex models, e.g. for large-capacity neural networks.

Two options are introduced in [2] for avoiding this problem. The first approach, instead of cross-entropy loss, suggests using mean square regression directly on weights, taking advantage of the property (5):

$$L_2(f) = \sum_i (f(x_i) - w_i)^2. \quad (7)$$

and restricting values of $f(x)$ to $[0, 1]$, e.g. by representing it as $\sigma(g(x))$ where $\sigma(\cdot)$ is the sigmoid function, $g(x)$ is an unbounded regression model.

Another approach is to use the maximum likelihood principle and avoid sPlot altogether and use the known probabilities $P(S | m)$ directly:

$$L_3(f) = - \sum_i \log [f(x_i)p(m_i | S) + (1 - f(x_i))p(m_i | B)]. \quad (8)$$

A detailed analysis of these loss functions and formal proofs can be found in [2].

3. Our approaches

In this work we extend approaches discussed above to the case of classification between two classes (C_1 and C_2), both of them represented as mixtures of signal (denoted as S_1 and S_2) and background (denoted as B_1 and B_2) events: $\{(x_i, m_i, y_i)\}_{i=1}^N$, where x_i are samples of target variables, m_i samples of discriminative variables, and y_i is binary class indicator.

Consider cross-entropy loss between S_1 and S_2 :

$$\mathcal{L}_4(f) = P(S_1 | S_1 \cup S_2) \mathbb{E}_{x \sim S_1} \log f(x) + P(S_2 | S_1 \cup S_2) \mathbb{E}_{x \sim S_2} \log(1 - f(x)) \quad (9)$$

where $\mathbb{E}_{x \sim S} h(x)$ denotes the conditional expectation $E_x[h(x) | S]$, $P(S | S_1 \cup S_2)$ — the probability of S given S_1 or S_2 .

$$\begin{aligned} \mathcal{L}_4(f) &= \alpha \mathbb{E}_{x \sim S_1} \log f(x) + \beta \mathbb{E}_{x \sim S_2} \log(1 - f(x)) = \\ &= \alpha \mathbb{E}_{x \sim C_1} \frac{P(x | S_1)}{P(x | C_1)} \log f(x) + \beta \mathbb{E}_{x \sim C_2} \frac{P(x | S_1)}{P(x | C_1)} \log(1 - f(x)) = \\ &= \alpha \mathbb{E}_{x \sim C_1} \frac{P(S_1 | x, C_1)P(x | C_1)}{P(S_1 | C_1)P(x | C_1)} \log f(x) + \beta \mathbb{E}_{x \sim C_2} \frac{P(S_2 | x, C_2)P(x | C_2)}{P(S_2 | C_2)P(x | C_2)} \log(1 - f(x)) = \\ &= \frac{\alpha}{P(S_1 | C_1)} \mathbb{E}_{x \sim C_1} P(S_1 | x, C_1) \log f(x) + \frac{\beta}{P(S_2 | C_2)} \mathbb{E}_{x \sim C_2} P(S_2 | x, C_2) \log(1 - f(x)); \quad (10) \end{aligned}$$

where: $\alpha = P(S_1 | S_1 \cup S_2)$, $\beta = P(S_2 | S_1 \cup S_2)$. Therefore, loss (10) can be approximated by:

$$L_4(f) = \sum_i \frac{y_i p_1(x_i)}{c_1} \log f(x_i) + \frac{(1 - y_i) p_2(x_i)}{c_2} \log(1 - f(x_i)); \quad (11)$$

where c_1 and c_2 are estimates of $P(S_1 | C_1)$ and $P(S_2 | C_2)$ correspondingly, $p_{1,i}$ and $p_{2,i}$ are estimates of $P(S_1 | x, C_1)$ and $P(S_2 | x, C_2)$ obtained by e.g. exact likelihood approach from [2].

Alternatively, with the help of sPlot weights, loss (9) can be expressed through binary indicator y :

$$\mathcal{L}_4(f) = \mathbb{E}_{x,y,m} y \cdot w_1(m) \log f(x) + (1 - y) \cdot w_2(m) \log(1 - f(x)), \quad (12)$$

where $w_1(m)$ and $w_2(m)$ denote weights assigned by sPlot for classes C_1 and C_2 correspondingly.

Here we also encounter potentially negative weights and all the problems associated with them as discussed in section 2. Applying weights averaging as in [2], one can avoid negative weights:

$$\mathcal{L}_4(f) = \mathbb{E}_{x,y} \left(\mathbb{E}_m [w_1(m) | x] \right) y \log f(x) + \left(\mathbb{E}_m [w_1(m) | x] \right) (1 - y) \log(1 - f(x)) \approx \sum_{i=1}^N y_i p_1(x_i) \log(f(x_i)) + (1 - y_i) p_2(x_i) \log(1 - f(x_i)), \quad (13)$$

where $p_1(x)$, $p_2(x)$ can be estimated by a regression trained on the corresponding weights with the onstrained MSE loss (7).

4. Experimental evaluation

We evaluate the proposed methods on data collected from the LHCb Muon subsystem [3] (MuID). To our knowledge, this is the only open dataset [4] that has sWeights. The dataset contains different particle species obtained from different calibration decays [5] and is labeled. Each species has its own background that is subtracted with a separate application of the sPlot method. For simplicity, we use only pion and muon species and ignore the kinematic weights. We use 60 scalar features. We split the dataset into train and test parts containing $2 \cdot 10^6$ and $6 \cdot 10^5$ examples respectively. For each train dataset size, the classifier was fitted 10 times to estimate the standard deviation of the scores.

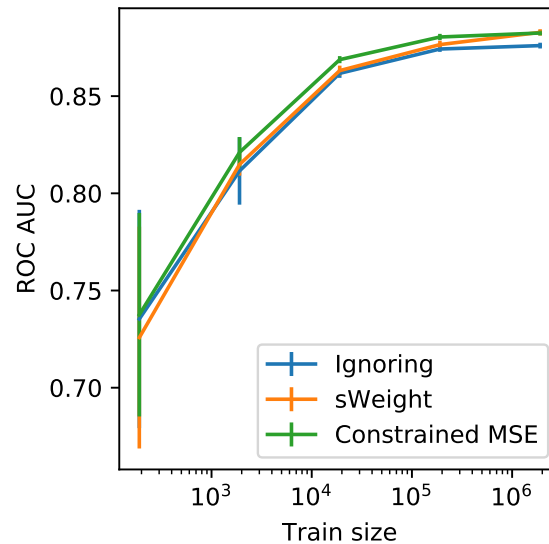


Figure 2: Experimental evaluation of CatBoost with the presented loss on the MuID dataset. sWeights – using sWeights directly as weights for logloss; Constrained MSE – our regression (7) used to transform the sWeights, the result used as weights for classification; Ignoring weight – training ignoring the sWeights.

The results are on the figure 2. The models used are described in Appendix A. Direct application of sWeights statistically significantly loses to our approach for some training sizes, performing about the same for the rest. The gap decreases with the training set size increase.

However, ignoring the sWeights during training also yields good results, so it seems that the impact of sWeights on the problem is limited. This suggests, that the distribution of pionic background has good separation from the distribution of muon signal and vice versa. In this case, a classifier that is optimal for separating the muon signal/background mixture from the pion signal/background mixture is also almost optimal for the separation of muon signal from the pion signal. This limits the utility of the dataset with the respect to evaluating the proposed methods.

To address this problem, we do an additional test with synthetically constructed sWeights, that are tailored to impact muon/pion classification. There are many equally attractive ways to introduce such sWeights. We go with the most obvious one: mark a sample of particles that are muons with a high level of confidence as pion background. First, we train a CatBoost classifier to separate signal and background and a classifier to separate muons and pions as described in section 2. We select 30% of examples with muon label with the highest scores and marked them as pion background. Next, we use the signal probabilities obtained from the first step and assign definite signal and background labels by cutting on its output. We generate “pseudomass” for both muon and pion background, each from a different distribution, which we use to compute the new sWeights.

As the result, we increase the impact of the sWeights on the classification problem to the point, where ignoring the sWeights leads to a significant drop in performance, while both our method and the baseline (using sWeights as example weights) are affected equally to permit for a fair comparison. The results are present in figure 3. Surprisingly, for CatBoost with train size equal to $2 \cdot 10^3$, using sWeights as example weights has the best performance. In all other cases, our methods perform better than using sWeights as example weights during classifier training.

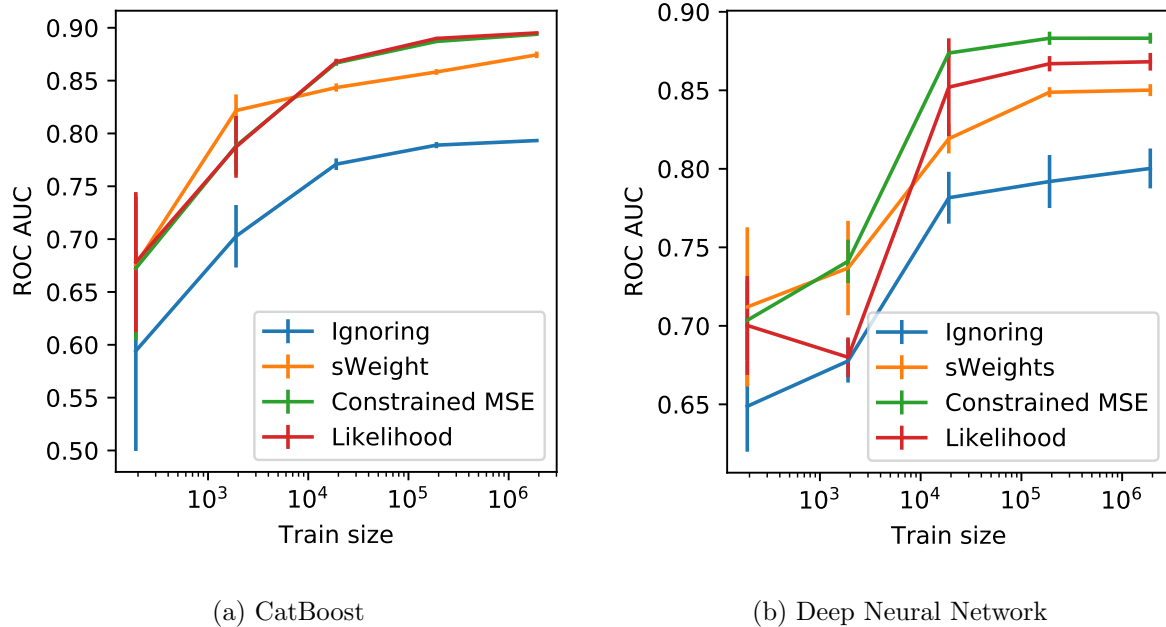


Figure 3: Experimental evaluation of our loss functions on the MuID dataset with artificial sWeights. sWeights – using sWeights directly as weights for logloss; Constrained MSE – our regression (7) used to transform the sWeights, the result used as weights for classification; Likelihood – our likelihood (8) for the transformation; Ignoring weight – training ignoring the sWeights.

5. Conclusion

Training a machine learning algorithm on mixtures of signal and background events in the absence of exact labels is a challenging task. With discriminative variables available, the sPlot technique becomes a tempting solution. However, it might potentially lead to loss function with no lower bound, which in turn might result in catastrophic overfitting.

In this work, we consider the problem of separating two classes, each contaminated by background events. Building upon the prior work on signal-background separation, we introduce two loss functions that guarantee the stability of training by avoiding negative weights. Experimental evaluation shows improved classification quality compared to using sWeights directly as example weights, both for neural network and gradient boosting model

6. Acknowledgments

The research leading to these results has received funding from Russian Science Foundation under grant agreement n 17-72-20127.

We are grateful to the LHCb collaboration for opening their data; Artem Maevskiy for the weighted ROC AUC code.

7. References

- [1] Pivk M and Le Diberder F R 2005 *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **555** 356–369
- [2] Borisyak M and Kazeev N 2019 *Journal of Instrumentation* **14** P08020–P08020 URL <https://doi.org/10.1088/1748-0221/14/P08020>
- [3] Archilli F, Baldini W, Bencivenni G, Bondar N, Bonivento W, Cadeddu S, Campana P, Cardini A, Ciambrone P, Vidal X C *et al.* 2013 *Journal of Instrumentation* **8** P10020
- [4] LHCb collaboration 2019 Muon identification challenge at IDAO URL <https://idao.world/>
- [5] Aaij R, Anderlini L, Benson S, Cattaneo M, Charpentier P, Clemencic M, Falabella A, Ferrari F, Fontana M, Gligorov V V *et al.* 2019 *EPJ Techniques and Instrumentation* **6** 1
- [6] Kingma D P and Ba J 2014 *arXiv preprint arXiv:1412.6980*

Appendix A. Models parameters

- Fully-connected neural networks (NN): 4 layers, 512, 256, 128 neurons in the first three layers and either 1 or 2 neurons in the last one, leaky ReLu (0.05), optimized by adam [6] algorithm with learning rate $2 \cdot 10^{-4}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, trained for 32 epochs; regressors and classifiers use the same architecture;
- CatBoost: 500 trees, leaf_estimation_method="Gradient", version 0.10.2 with our losses added and check for negative weights removed: https://github.com/kazeevn/catboost/tree/constrained_regression