RECEIVED: August 6, 2020

REVISED: January 11, 2021

ACCEPTED: February 20, 2021

PUBLISHED: March 31, 2021

ThickBrick: optimal event selection and categorization in high energy physics. Part I. Signal discovery

Konstantin T. Matchev and Prasanth Shyamsundar

*Institute for Fundamental Theory, Physics Department, University of Florida,
Gainesville, FL 32611, U.S.A.*

E-mail: matchev@ufl.edu, prasanths@ufl.edu

ABSTRACT: We provide a prescription called ThickBrick to train optimal machine-learning-based event selectors and categorizers that maximize the statistical significance of a potential signal excess in high energy physics (HEP) experiments, as quantified by any of six different performance measures. For analyses where the signal search is performed in the distribution of some event variables, our prescription ensures that only the information complementary to those event variables is used in event selection and categorization. This eliminates a major misalignment with the physics goals of the analysis (maximizing the significance of an excess) that exists in the training of typical ML-based event selectors and categorizers. In addition, this decorrelation of event selectors from the relevant event variables prevents the background distribution from becoming peaked in the signal region as a result of event selection, thereby ameliorating the challenges imposed on signal searches by systematic uncertainties. Our event selectors (categorizers) use the output of machine-learning-based classifiers as input and apply optimal selection cutoffs (categorization thresholds) that are functions of the event variables being analyzed, as opposed to flat cutoffs (thresholds). These optimal cutoffs and thresholds are learned iteratively, using a novel approach with connections to Lloyd's k-means clustering algorithm. We provide a public, Python implementation of our prescription, also called **ThickBrick**, along with usage examples.

KEYWORDS: Supersymmetry Phenomenology

ARXIV EPRINT: [1911.12299](https://arxiv.org/abs/1911.12299)

Contents

1	Introduction	1
1.1	Dimensionality reduction and complementary information	3
1.2	The landscape of related methods	6
2	Groundwork	8
2.1	Setup and notation	8
2.2	Statistical distances	9
2.3	Eventwise contribution to $D_{\text{Neym}\chi^2}$	12
2.4	Iterative maximization of $D_{\text{Neym}\chi^2}$	14
2.5	Generalizing to other statistical distances	16
2.6	The rectified “misalignment”	17
3	Prescription for training an optimal event categorizer	18
3.1	Preprocessing: learning $p(e)$	19
3.2	Choices to be made	21
3.3	Iterative training prescription	22
3.4	Performance evaluation	24
3.5	Miscellaneous notes	25
4	The ThickBrick package: a public implementation in Python	26
5	Conclusions	33
6	Epilogue: decorrelated classifiers and taggers	34
A	Statistical distances	36
A.1	$D_{\text{Neym}\chi^2}$ and $D_{\text{Pear}\chi^2}$: asymptotic expected values of Neyman- χ^2 and Pearson- χ^2 statistics	37
A.2	D_{KL} : Kullback-Leibler divergence from null to alternative	38
A.3	D_{revKL} : Kullback–Leibler divergence from alternative to null	39
A.4	D_{J} : Jeffreys divergence	39
A.5	D_{B} : Bhattacharyya distance	39
B	An interpretation of Kullback-Leibler divergence	40
C	Proof of correctness of the iterative training algorithm	43
D	A comparison of the different distance measures	45

1 Introduction

The increase in complexity and scale of high energy physics (HEP) experiments over the years has been accompanied by an increased sophistication of the techniques employed in the data analysis. Among the more recent additions to the medley of analysis techniques used in experimental HEP are machine learning algorithms. Machine learning techniques have found applications in various aspects of the analysis [1–3], and are an active area of research, both from the standpoint of finding new applications as well as refining their current usage.

An important step in any collider¹ data analysis is event selection and/or categorization. Event selection is the process of selecting a subset of collider events for further analysis. Event categorization is the process of grouping the selected events into disjoint categories, each of which will be individually analyzed.

Event selection serves two purposes. First, it helps bring down the volume of recorded data within the computational limits of the experiment (storage and processing power). This is done at the trigger level [4, 5] and works by making rapid decisions on whether or not an event is “interesting enough” to warrant further processing, given the goals of the experiment and its computational limitations. Secondly, event selection is used in individual analyses to reduce the amount of “background” in the event sample and make it more “signal”-rich, with signal and background events defined appropriately. This is the type of event selection we will focus on in this paper.

A sample of collider events can often be thought of as being composed of several subsamples, although real collider events do not necessarily carry a label indicating which subsample they are from. These *mixture models* are used in the Monte Carlo simulation of events, with simulated events carrying a label to identify the subsample they are from. These subsamples can be labeled as signal and background (or bg1, bg2, etc), as appropriate for the particular analysis at hand. The event selection process can then be thought of as selectively including signal-like events (and rejecting background-like events) to “improve” the analysis. Let us call this the “signal is better than background” heuristic.

The goal of event selection in an analysis searching for a signal (due to new physics or an interesting Standard Model process), is to increase the expected significance of a potential signal excess, or improve the expected upper-limit on the signal cross-section (in the case of a null result). It is easy to see how the heuristic mentioned above can help improve such an analysis. On the other hand, the goal in a parameter measurement analysis (say a top quark or a W -boson mass measurement) is to increase the precision of the measurement. In such analyses, subsamples whose distributions are independent of the parameter(s) being measured can be labeled as background, and subsamples whose distributions do depend on the parameter(s) being measured can be labeled as signal. Again, we can see how the “signal is better than background” heuristic can improve the precision of measurement by making the distribution of the selected event sample more sensitive to the parameter(s) being measured.

¹In what follows, for brevity we shall refer to “collider” events such as those collected at the Large Hadron Collider (LHC), but our treatment is equally applicable to other areas of HEP as well, and maybe even outside physics, provided they share similar goals and have access to properly labeled training data.

The “signal is better than background” heuristic allows us to rephrase the event selection problem as an event classification problem. Over the last two decades, machine learning methods like artificial neural networks [6] and boosted decision trees (BDT) [7, 8] originally developed to solve classification problems have been successfully employed in collider physics analyses [1–3]. These algorithms are trained on simulated data samples to learn features that distinguish signal and background events. They can then provide a measure of how signal-like an event in the real data is, and the event selection can be done based on this measure.

But despite the rationale behind the “signal is better than background” heuristic, it is not perfectly aligned with the physics goals of improving the search sensitivity and measurement precision. In this work, we will prioritize the physics goals, and focus on maximizing the expected sensitivity of the analysis that uses the data after event selection/categorization. We will introduce appropriate performance measures that capture this sensitivity and provide prescriptions to maximize them. This is the optimality alluded to in the title of the paper, and it is in this sense that the word “optimal” is used throughout.

In this work, we introduce a paradigm named ThickBrick to construct optimal event selection and categorization prescriptions for physics analyses within the supervised machine learning paradigm, with cost functions defined on a per-event basis. Along the way, we will point out exactly how our methods realign the goals of event selection and categorization with the physics goals mentioned above. To facilitate clarity of presentation and improve accessibility, we will present the ThickBrick approach in three parts (tentative).

- In part I (this paper), we will focus on event selection and categorization for signal discovery (and upper-limit setting). We will provide a prescription to train event selectors and categorizers, which are optimal for an exactly specified signal distribution.
 - We will not account for unspecified signal parameter(s), like the mass of the dark photon in a dark photon search. In other words, in this paper we will train our categorizers to be optimal for a given set of values of signal parameters. We will also not concern ourselves with systematic uncertainties in the model specification. Both of these issues will be addressed in part III.
 - The technique introduced in part I is directly applicable in searches where the signal parameters are known beforehand, e.g., rare Higgs signatures, where the mass of the Higgs is known a priori.
- In part II [9], we will provide event selection and categorization prescriptions for reducing the expected *statistical* uncertainty of the measurement in parameter measurement experiments.
 - As in part I, we will not be accounting for model (systematic) uncertainties, which will be addressed in part III.
- In part III [10], we will provide prescriptions to train optimal event selectors and categorizers accounting for unspecified signal parameters and systematic uncertainties. In other words,

- We will provide prescriptions to train event selectors and categorizers that improve sensitivity over a range of unspecified signal parameter values.
- We will also provide prescriptions to maximize the significance of a signal excess in the presence of model uncertainties, and minimize the combined statistical and systematic uncertainty in parameter measurement experiments.

1.1 Dimensionality reduction and complementary information

Data recorded in collider experiments is typically high dimensional. For example, in LHC experiments, a typical recorded event leaves thousands of hits in the detector. Even at the parton level, depending on the number of reconstructed final state particles, events can have ~ 10 or more attributes. Ideally one would want to draw statistical inferences on the underlying physics by analyzing the distribution of events in these high dimensional spaces.

But analyzing the distribution of collider data in a high dimensional space (using multi-dimensional histograms, likelihood-ratio tests, the Matrix Element Method [11–13], or machine learning-based inference [14–17]²) comes with several challenges — the computational power needed to systematically scan the phase-space grows exponentially with the number of dimensions; the volume of data (real and simulated) needed to populate the space grows exponentially with the number of dimensions; directly evaluating the likelihood-ratio using the Matrix Element Method is computationally expensive [13], Monte Carlo validation in high dimensional phase-space (e.g., in the context of machine learning) can be challenging, etc.

One way to deal with the curse of dimensionality is to reduce the dimensionality — by constructing event variables like pseudorapidities, transverse momenta, invariant or transverse masses of sets of final state particles, etc. After reducing the dimensionality of data to a more tractable number of event variables, we can limit the analyses to studying the distribution of these event variables.

But this dimensionality reduction is accompanied by information loss as expected from the data processing inequality [19] — typically the dimensionality-reduced data contains less information than the original data, be it for the purpose of signal search or parameter measurement. In the case of signal search analyses, the information loss can be understood as follows: all the information contained in a collider event for the purpose of testing for the presence of a signal over a background can be captured by the likelihood-ratio of the event under the signal and background distributions (a consequence of the Neyman-Pearson lemma [20]). Reducing the dimensionality of the data mixes up regions of phase-space with different values of this likelihood-ratio, but the same value of the event variable(s). To see how such a mixing causes information loss, consider the simple scenario with two regions of phase-space where the expected number of signal and background events are, respectively, (S_1, B_1) and (S_2, B_2) , all non-zero. Let the events in the two regions be mixed and analyzed together, and let the combined (expected) signal and background counts be $(S_{\text{tot}}, B_{\text{tot}}) \equiv (S_1 + S_2, B_1 + B_2)$. Note that

$$B_1 \frac{S_1}{B_1} + B_2 \frac{S_2}{B_2} = (B_1 + B_2) \frac{S_{\text{tot}}}{B_{\text{tot}}} . \quad (1.1)$$

²For a comparison between ML-based approaches and the Matrix Element method see ref. [18].

This implies, by Jensen’s inequality [19, 21], that for any convex function f_{convex} ,

$$B_1 f_{\text{convex}}\left(\frac{S_1}{B_1}\right) + B_2 f_{\text{convex}}\left(\frac{S_2}{B_2}\right) \geq B_{\text{tot}} f_{\text{convex}}\left(\frac{S_{\text{tot}}}{B_{\text{tot}}}\right) , \quad (1.2)$$

with equality occurring iff $S_1/B_1 = S_2/B_2$. By choosing appropriate forms of f_{convex} , it can be shown that

$$\frac{S_1^2}{B_1} + \frac{S_2^2}{B_2} \geq \frac{S_{\text{tot}}^2}{B_{\text{tot}}} , \quad (1.3a)$$

$$\frac{S_1^2}{N_1} + \frac{S_2^2}{N_2} \geq \frac{S_{\text{tot}}^2}{N_{\text{tot}}} , \quad (1.3b)$$

$$\left[-S_1 + N_1 \ln\left(\frac{N_1}{B_1}\right) \right] + \left[-S_2 + N_2 \ln\left(\frac{N_2}{B_2}\right) \right] \geq -S_{\text{tot}} + N_{\text{tot}} \ln\left(\frac{N_{\text{tot}}}{B_{\text{tot}}}\right) , \quad (1.3c)$$

where

$$N_1 \equiv S_1 + B_1 , \quad N_2 \equiv S_2 + B_2 , \quad N_{\text{tot}} \equiv S_{\text{tot}} + B_{\text{tot}} .$$

Eq. (1.3) illustrates the loss of sensitivity of the analysis to the presence of a signal caused by phase-space mixing or dimensionality reduction, with sensitivity quantified by several measures, including the familiar “ $S/\sqrt{B}$ or $S/\sqrt{N}$ added in quadrature”.

Many commonly used analysis techniques, including event selection and categorization, can be thought of as attempts to mitigate the information loss caused by dimensionality reduction. Event categorization helps by separating regions of phase-space, *that would otherwise be mixed together*, into disjoint categories. Similarly, event selection would help when not including certain parts of phase-space leads to better sensitivity than mixing them with regions of higher purity — in our example, if $S_1^2/B_1 > S_{\text{tot}}^2/B_{\text{tot}}$, then rejecting region 2 is better than mixing the two regions together.

Let $\mathbf{e}$ represent all the observable information pertaining to a single collider event. Let $\mathbf{x}$ be an event variable constructed out of $\mathbf{e}$ whose distribution will be analyzed to search for the presence of a signal. In order to mitigate the loss of information caused by dimensionality reduction $\mathbf{e} \rightarrow \mathbf{x}$, event selection or categorization needs to be done based on the information lost during dimensionality reduction, i.e., information complementary to $\mathbf{x}$. Event selection or categorization based on the information contained in $\mathbf{x}$ itself cannot improve the sensitivity of the experiment, and can potentially be detrimental to it.³ This is illustrated in figure 1, which shows event selection performed based only on $\mathbf{x}$. This only removes some bins from the analysis, hence removing some non-negative, potentially positive, terms from the sum $\sum_{i \in \mathbf{x} \text{ bins}} s_i^2/b_i$ (or other measures of sensitivity).

Typically event selection based on machine learning (ML) classifiers (trained on the full event information $\mathbf{e}$) is done by choosing a working point or selection threshold on the ML output. The problem here is that the ML classifier learns to distinguish between signal and background based on the full information in $\mathbf{e}$, and not just the complementary information. As a result, this selection procedure has to find a compromise between the

³Even in the absence of dimensionality reduction, event selection and categorization can still be useful in computationally expensive analyses (like the Matrix Element Method) to reduce the size of the dataset, and hence the computational cost, with an acceptably low loss of sensitivity.

Figure 1. An illustration of event selection based solely on the event variable to be analyzed after the selection process. The dark shaded region represents the rejected region of phase-space. This kind of event selection does not improve the sensitivity of the analysis. To improve sensitivity, background needs to be selectively removed from each bin as opposed to removing entire bins.

gain in sensitivity from tapping into complementary information and the loss in sensitivity resulting from using non-complementary information. This problem is particularly severe when $\mathbf{x}$ is a “good event variable” in terms of separating signal and background, since such a variable will be strongly correlated with the classifier output.

This is the source of “misalignment” addressed in this paper. As we will demonstrate, such a compromise is not necessary, and it is possible to force an event selector to learn from only information complementary to $\mathbf{x}$. This can be done by using different selection thresholds at different values of $\mathbf{x}$. Intuitively this can be understood as follows: to maximize $\sum_{i \in \mathbf{x} \text{ bins}} s_i^2/b_i$ using an event selector, one needs to maximize s_i^2/b_i in each $\mathbf{x}$ -bin, by choosing an appropriate, bin dependent, threshold on the ML classifier output. On the other hand, even the best bin independent threshold will only be a compromise value.

In the case of event categorizers, although using information contained in the event variable $\mathbf{x}$ is not detrimental to the sensitivity, it still does not improve it. So a similar compromise between using complementary and non-complementary information exists in typical ML classifier based categorization as well (flat thresholds between categories), which can also be avoided using $\mathbf{x}$ dependent thresholds between categories.

In addition to leading to sub-optimal event selectors and categorizers in terms of statistical sensitivity, non-complementary information used in event selection often causes the background distribution (post-selection) to be peaked in the same region as the signal. This amplifies the challenge posed to signal searches by systematic uncertainties in the background specification. Although this paper is not geared towards reducing systematic uncertainties (which we will do in part III [10]), event selectors decorrelated from the event variables $\mathbf{x}$ can make it easier to estimate the background distribution from data itself, thereby reducing the impact of systematic uncertainties. This is illustrated in figure 2 using plots from a toy example discussed in section 4. The left panel shows the distributions of signal and background before selection. The remaining two panels show their distribution post-selection using event selectors based on the full information $\mathbf{e}$ (middle panel)

Figure 2. Plots from the toy example considered in section 4. The left panel shows the distribution of signal and background before event selection, with the background normalized to 1 and a signal-to-background ratio of 1:9. The middle panel shows the result from a selection based on full event information, including the event variable x . The right panel shows the result from the procedure outlined in section 3 below, which uses only information complementary to the event variable x .

and complementary information only (right panel). We observe that in the middle panel, the background distribution becomes sculpted and peaks in the same region as the signal, while in the right panel, with the procedure which we advocate below, the background distribution remains flat.

1.2 The landscape of related methods

It was pointed out in [22, 23] that the performance of event selectors in terms of classification accuracy is not perfectly correlated with the sensitivity and precision of the physics analysis that uses them — neural networks trained to directly optimize for sensitivity (as in a Higgs search) or precision (as in a top mass measurement) were demonstrated to outperform event selectors trained for maximal classification accuracy (with thresholds or working points chosen to optimize precision). It is precisely this misalignment that we seek to explain and alleviate in this paper and the subsequent⁴ work [9, 10].

In [22, 23], an evolutionary method for training neural networks called NeuroEvolution of Augmenting Topologies (NEAT) [24] was used in the direct optimization of event selectors. Another approach to directly optimizing neural networks for discovery significance was introduced in [25], using cost functions defined over batches of training events. Thick-Brick, on the other hand, relies on simply using the output of typical ML-based classifiers in an optimal manner.

At the heart of this paper lies the following observation: the task of finding the event selector/categorizer that maximizes the statistical sensitivity of a search analysis (performed in the distribution of an event variable) is intrinsically related to the space of event selectors/categorizers over which the maximization is performed. We claim that selectors/categorizers based only on classifier output (flat thresholds) lack the flexibility needed to achieve global optimality (i.e., in the space of **all** selectors/categorizers), and expanding the space of categorizers under consideration can potentially lead to improved results. This is in line with the findings of [26], where, in addition to using thresholds on a BDT classifier output, the flexibility of cutting on another carefully chosen variable was added to the categorization process, leading to a better performing categorizer. We show in sec-

⁴The findings in [22, 23] relate more to the forthcoming parts 2 and 3 than the current paper.

tion 2 that, from an information theory standpoint, global optimality can be attained with our proposal to use the ML classifier output as well as the event variable $\mathbf{x}$, and impose $\mathbf{x}$ -dependent thresholds.

Our work is closely related to refs. [27, 28] which consider the optimal partitioning of events into different categories, in order to maximize the sensitivity of an analysis that searches for a signal or measures a parameter by counting the number of events in each category. They approach this task by maximizing the Fisher information [19] left in the data after reducing each event to the category it belongs to. Parts 1 and 2 of our work can be viewed as maximizing the Fisher information left in the data after reducing each event to the category it belongs to *and an event variable $\mathbf{x}$* . Note that a signal search problem can be viewed as a signal cross-section measurement problem, and the performance measures introduced in our paper reduce to the relevant Fisher information in the faint signal limit (when the ‘background only’ and ‘background+signal’ hypotheses are only infinitesimally apart).

Refs. [29, 30] considered double Higgs production in the $b\bar{b}\gamma\gamma$ channel, and performed a simultaneous optimization of cuts on kinematic variables and BDT hyperparameters, for an analysis that used both filters before counting the number of events that pass them. For such a counting analysis (as opposed to a distribution analysis), applying a threshold on the score from a perfectly trained classifier should yield the highest sensitivity (from the Neyman-Pearson lemma). However, due to practical limitations the BDT training cannot be perfect, and cuts on kinematic variables can be used to improve the performance.

Several approaches exist in the literature for decorrelating the output of a classifier from a particular event variable. One possibility is to use only features decorrelated from the event variable as inputs to the ML classifier [31–33]. Another way is to force decorrelation during training with appropriate penalties [34–36]. In our case, the decorrelation is achieved naturally as a byproduct of maximizing performance measures more in line with the sensitivity of the physics analysis than, say, classification accuracy. Another technique related to the correlation between variables is the *sPlot* tool [37] which can extract the distribution of a variable x within each subsample of events, given a mixture of the different subsamples and a variable y , uncorrelated with x , whose distribution within each subsample is known.

The rest of this paper is organized as follows: in section 2, we will introduce several performance measures for event selectors/categorizers, and lay the groundwork for a training prescription to maximize them. We will also explain in more detail how our methods realign the goals of event selection and categorization with the physics goals. Section 3 will be a description of our training algorithm, intended to serve as an implementation and usage guide. To this end, we will keep the material in section 3 mostly self-contained, with minimal dependence on section 2. In section 4, we will demonstrate the working of our algorithm using a toy example analyzed using the [ThickBrick package](#) [38], our publicly available implementation of the algorithm. Finally, we will provide some outlook on the future of this work and possible extensions in sections 5 and 6.

2 Groundwork

2.1 Setup and notation

In this subsection we will establish the setup and notation for a generic signal search procedure considered in this paper. Throughout this paper, we will work under the frequentist approach to statistical inference,⁵ so no notion of probability will be associated with the presence or absence of a signal.

Let $\mathcal{S}$ be a collection of collider events. We will treat a given instance of $\mathcal{S}$ as one of the possible outcomes of the experiment performed to produce it. The events in $\mathcal{S}$ are assumed to be independent and identically distributed. The total number of events in $\mathcal{S}$ is assumed to be a Poisson distributed random variable whose mean could possibly depend on the underlying theory. In practice, $\mathcal{S}$ can be a collection of events recorded at the collider that pass some event pre-selection criteria.

Let $\mathbf{e}$ be a random vector which contains all the observable information (that is available to an analysis) pertaining to a single collider event. In practice, $\mathbf{e}$ can include *estimated* number and ids of parton-level final state particles in the event, their *reconstructed* (i.e., after detector smearing) energies and momenta, etc.

Let $\mathbf{x}$ be an event variable (also a random vector) constructed out of $\mathbf{e}$ whose distribution is to be analyzed to search for the presence of a signal. For example, $\mathbf{x}$ could be the 4-lepton invariant mass for a Higgs search in the 4-lepton channel.

In general, the vectors $\mathbf{e}$ and $\mathbf{x}$ can have continuous (real valued) and discrete components.⁶ Let $\Omega_{\mathbf{e}}$ and $\Omega_{\mathbf{x}}$ represent, respectively, the sample spaces $\mathbf{e}$ and $\mathbf{x}$ belong to, with appropriate probability measure functions.⁷

We consider a generic signal search analysis performed in the following two steps:

Step 1. Event selection/categorization. From $\mathcal{S}$, events are categorized based on their $\mathbf{e}$ values into C categories ($C \geq 1$) numbered $1, \dots, C$. If we allow for event rejection, an event can also be not included in any of the C categories (rejected) based on its $\mathbf{e}$ value. For now we will assume that the number of categories, C , is predetermined.

Let $\mathcal{S}_c$ be the collection of events in category c for $c \in \{1, \dots, C\}$.

Note that if we use only one category and allow for event rejection, this behaves as a simple in-or-out event selection procedure. *Henceforth, “event categorization” will refer to both event selection and event categorization, unless a distinction is explicitly made.*

⁵The techniques we will introduce can have applications in Bayesian inference as well. Under Bayesian inference, the goal of an analysis would be to reduce the expected posterior entropy of the presence or absence of a signal (i.e., reduce the uncertainty on whether or not the signal exists). Our techniques can *roughly* be interpreted as minimizing this expected entropy.

⁶For categorical/discrete valued components of $\mathbf{e}$, like the number and ids of reconstructed final state particles, it is a good idea in practice to group events based on those features and analyze each group separately, although our formalism does not assume this to have been done.

⁷Using as σ -algebra, the product of Borel algebras for real valued components and powersets for discrete valued components.

Step 2. Number density analysis. The event variable $\mathbf{x}$ is calculated for events in each of the C categories.

The number density distribution of $\mathbf{x}$ in each of the categories $\mathcal{S}_c$ is then analyzed to test for the presence of a signal, say, using the likelihood-ratio test or the χ^2 difference test.

We will assume that there are two subpopulations of events, background and signal, each with its own exactly specified *number density* distribution of $\mathbf{e}$. The null hypothesis, $\mathbb{H}_b$, is that only the background subpopulation exists in the true event population, and the alternative hypothesis, $\mathbb{H}_{b+s}$, is that both signal and background events exist in the true event population.

The goal of this paper is to provide a prescription to perform the event categorization step (step 1) “optimally”. To this end, we need a measure of how distinguishable the two hypotheses are after reducing each event to its $\mathbf{x}$ value and the category it belongs to, i.e., after event categorization. We lay the notational groundwork for that in table 1.

Note that both the null and the alternative hypotheses are point hypotheses, i.e., there are no free parameters in either $\mathbb{H}_b$ or $\mathbb{H}_{b+s}$, including the relative total abundance of signal and background, S/B . The effect of scaling the signal and background cross-sections relative to each other⁸ will become apparent at the end of the construction.

2.2 Statistical distances

To formalize the problem of optimizing the event selection/categorization procedure, we need a measure of how distinguishable the null and the alternative hypotheses are after reducing each event to the category it belongs to and its $\mathbf{x}$ value. In eq. (2.1) we provide six such measures, referred to as *statistical distances* between the hypotheses. The prescription we develop in this paper is compatible with any of these distance measures; they represent a choice available when using ThickBrick.

$$D_{\text{Neym}\chi^2} = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \frac{s_c^2(\mathbf{x})}{n_c(\mathbf{x})} \quad (2.1a)$$

$$D_{\text{Pear}\chi^2} = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \frac{s_c^2(\mathbf{x})}{b_c(\mathbf{x})} \quad (2.1b)$$

$$D_{\text{KL}} = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \left[-s_c(\mathbf{x}) - n_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] \quad (2.1c)$$

$$D_{\text{revKL}} = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \left[s_c(\mathbf{x}) + b_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] \quad (2.1d)$$

$$D_J = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \left[-s_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] \quad (2.1e)$$

$$D_B = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \left[n_c(\mathbf{x}) - \frac{s_c(\mathbf{x})}{2} - n_c(\mathbf{x}) \sqrt{1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})}} \right] \quad (2.1f)$$

⁸An overall scaling of signal and background number densities (keeping S/B fixed) will have no effect on our event selection/categorization prescriptions.

Notation	Definition/Interpretation
$\mathbf{e}$	Random vector which contains all the observable information (that is available to an analysis) pertaining to a single collider event. Randomness of $\mathbf{e}$ refers to the randomness of collider events.
$\mathbf{x}$ or $\mathbf{x}(\mathbf{e})$	Event variable (also a random vector) constructed out of $\mathbf{e}$ whose distribution is to be analyzed to search for the presence of a signal.
$\Omega_{\mathbf{e}}$	The sample space of event $\mathbf{e}$.
$\Omega_{\mathbf{x}}$	The sample space of event variable $\mathbf{x}$.
$\mathbb{H}_b$	Null hypothesis that only the background subpopulation exists in the true event population.
$\mathbb{H}_{b+s}$	Alternative hypothesis that both background and signal subpopulations exist in the true event population.
$\eta(\mathbf{e})$	Represents the event categorizer used in the analysis. η is a function of $\mathbf{e}$ and returns the category (integer) between 1 and C to place an event in. η can also return None if event rejection is allowed.
B	Expected total number of events in the collection $\mathcal{S}$, under the null hypothesis $\mathbb{H}_b$. It is also the expected number of background events under $\mathbb{H}_{b+s}$.
N	Expected total number of events in the collection $\mathcal{S}$, under the alternative hypothesis $\mathbb{H}_{b+s}$.
S	$S \equiv N - B$. Expected number of signal events under $\mathbb{H}_{b+s}$.
$b(\mathbf{e})$	<i>Number density</i> distribution of $\mathbf{e}$ under the null hypothesis. It is proportional to the probability density function for $\mathbf{e}$ under $\mathbb{H}_b$, and is normalized to B .
$n(\mathbf{e})$	Number density distribution of $\mathbf{e}$ under the alternative hypothesis $\mathbb{H}_{b+s}$ (normalized to N).
$s(\mathbf{e})$	$s(\mathbf{e}) \equiv n(\mathbf{e}) - b(\mathbf{e})$. Number density distribution of signal events under the alternative hypothesis. It is non-negative and normalized to S .
B_c	Expected number of events in category c under $\mathbb{H}_b$. This also equals the expected number of background events in category c under $\mathbb{H}_{b+s}$. If event rejection is not allowed, $\sum_c B_c = B$.
N_c	Expected number of events in category c under $\mathbb{H}_{b+s}$. If event rejection is not allowed, $\sum_c N_c = N$.
S_c	$S_c \equiv N_c - B_c$. Expected number of signal events in category c under $\mathbb{H}_{b+s}$. If event rejection is not allowed, $\sum_c S_c = S$.
$b_c(\mathbf{x})$	Number density distribution of $\mathbf{x}$ in category c under $\mathbb{H}_b$ (background distribution). It is normalized to B_c .
$n_c(\mathbf{x})$	Number density distribution of $\mathbf{x}$ in category c under $\mathbb{H}_{b+s}$. It is normalized to N_c .
$s_c(\mathbf{x})$	$s_c(\mathbf{x}) \equiv n_c(\mathbf{x}) - b_c(\mathbf{x})$. Number density distribution of $\mathbf{x}$ in category c for signal events. It is normalized to S_c .
$p(\mathbf{e})$	$p(\mathbf{e}) \equiv \frac{s(\mathbf{e})}{n(\mathbf{e})}$ The probability that an event sampled from the alternative hypothesis $\mathbb{H}_{b+s}$ is a signal event, conditional on its $\mathbf{e}$ value.
$p_c(\mathbf{x})$	$p_c(\mathbf{x}) \equiv \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})}$ The probability that an event sampled from the alternative hypothesis $\mathbb{H}_{b+s}$ is a signal event, conditional on its $\mathbf{x}$ value and the category c it belongs to. $p_c(\mathbf{x})$ can also be thought of as the expected value of $p(\mathbf{e})$ conditional on $(c, \mathbf{x})$.

Note: all quantities with the subscript c implicitly depend on the categorizer η used.

Table 1. Notation.

We will only provide a brief description of these distances here, and relegate the derivation of these formulas to appendix A. We will go over the estimation of these integrals in section 2.3, and provide a detailed comparison between the distance measures in appendix D.

- $D_{\text{Neym}\chi^2}$ is the asymptotic expected value of the Neyman- χ^2 statistic [39], sometimes referred to as the modified χ^2 statistic, in the fine binning limit under the alternative hypothesis. This may be familiar as “ $s/\sqrt{n}$ summed over bins in quadrature”.
- $D_{\text{Pear}\chi^2}$ is the asymptotic expected value of the Pearson- χ^2 statistic [40] in the fine binning limit under the alternative hypothesis. Again, this may be familiar as “ $s/\sqrt{b}$ summed over bins in quadrature”.
- D_{KL} is the Kullback-Leibler divergence [41] from the null hypothesis to the alternative hypothesis. Up to a factor of 1/2, this form may be familiar as the (asymptotic) expected significance of excess under the alternative hypothesis, squared [42]. We provide an alternative interpretation of D_{KL} in appendix B.
- D_{revKL} ⁹ is the Kullback-Leibler divergence [41] from the alternative hypothesis to the null hypothesis. We provide an interpretation of D_{revKL} in appendix B.
- D_{J} is Jeffreys divergence [43] between the null and alternative hypotheses. It is a symmetrized version of the Kullback-Leibler divergence defined as $D_{\text{J}} \equiv D_{\text{KL}} + D_{\text{revKL}}$.
- D_{B} is the Bhattacharyya distance [44] between the null and alternative hypotheses.

In the limit $s_c(\mathbf{x}) \ll b_c(\mathbf{x})$ for all c and almost all $\mathbf{x}$, we have

$$D_{\text{Neym}\chi^2} \approx D_{\text{Pear}\chi^2} \approx 2D_{\text{KL}} \approx 2D_{\text{revKL}} \approx D_{\text{J}} \approx 8D_{\text{B}} \quad (2.2)$$

up to leading order in $s_c(\mathbf{x})/n_c(\mathbf{x})$. Generically we will refer to these distances as D_{stat} .

$$D_{\text{stat}} \in \left\{ D_{\text{Neym}\chi^2}, D_{\text{Pear}\chi^2}, D_{\text{KL}}, D_{\text{revKL}}, D_{\text{J}}, D_{\text{B}} \right\}.$$

Some key points about these distances relevant to our discussion are:

1. *They capture distinguishability of hypotheses.* The measures given in eq. (2.1) do not depend on actual data. Instead they depend on the two hypotheses, $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$, for the distribution of data after event categorization. The greater the value of these distances, the more distinguishable the hypotheses.

Although not explicit from their form, it can be shown that $D_{\text{stat}} \geq 0$, with equality occurring iff $b_c(\mathbf{x}) = n_c(\mathbf{x})$ for all $c \in \{1, \dots, C\}$ and almost all $\mathbf{x} \in \Omega_{\mathbf{x}}$, i.e., iff the hypotheses $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$ are experimentally indiscernible from one another.

⁹ D_{revKL} is not standard notation. We intend for D_{revKL} to be read as “reverse Kullback-Leibler (divergence)” or “reverse KL (divergence)”.

2. *They depend on the event categorization criteria.* How events are categorized implicitly affects each of the distance measures. This allows us to phrase the problem of optimizing the event categorization procedure as maximizing D_{stat} by varying the event categorization criteria.
3. *They scale linearly with the integrated luminosity of the experiment.* Scaling the number of signal and background events proportionally, without changing their distributions or relative abundance, will scale D_{stat} by the same factor. As a result, we can write each of our statistical distances as N times the expected value of a per-event contribution defined appropriately. We will use this in constructing a training prescription to maximize our distance measures.
4. *They only depend on $\mathbf{x}$ and p values of events* that get placed in each category. In particular, they depend on $n_c(\mathbf{x})$ and $p_c(\mathbf{x})$ — this is akin to the significance of excess depending on the number of events and signal-background ratio. Here $p_c(\mathbf{x})$ is the expected value of $p(\mathbf{e})$ for events in category c with the given value of $\mathbf{x}(\mathbf{e})$ (assuming events are sampled from $\mathbb{H}_{b+s}$). Recall that $p(\mathbf{e})$ and $p_c(\mathbf{x})$ are the probabilities that an event sampled from $\mathbb{H}_{b+s}$ is a signal event, conditioned on $\mathbf{e}$ and $(c, \mathbf{x})$, respectively.

This means that there exists an optimal categorization criterion based only on the $\mathbf{x}$ and p values of events. This is to be expected from the Neyman-Pearson lemma, a consequence of which is that $p(\mathbf{e})$ contains all¹⁰ the information in an event relevant to distinguishing between the hypotheses $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$. The optimal categorization criterion will also take into account the $\mathbf{x}$ value of an event to make sure that only information complementary to $\mathbf{x}$ is used in the categorization.

2.3 Eventwise contribution to $D_{\text{Neym}\chi^2}$

Let us choose $D_{\text{Neym}\chi^2}$ as a demo distance measure to construct our prescription. For notational simplicity, wherever applicable, $\mathbf{x}$ will refer to $\mathbf{x}(\mathbf{e})$ of the particular event $\mathbf{e}$ under consideration.

Let us begin with a (training) set of events sampled from $\mathbb{H}_{b+s}$, i.e., both signal and background events. We want to train an event categorizer on this event sample to maximize $D_{\text{Neym}\chi^2}$. We will assume that the values of $\mathbf{x}$ and $p(\mathbf{e})$ are known for each event. $p(\mathbf{e})$ can be estimated *separately* using machine learning techniques as shown in [14, 17]. In section 3.1, we will discuss some practical aspects of learning $p(\mathbf{e})$ for our event categorization purpose.

In the interest of providing some intuition for what follows, let us momentarily sacrifice mathematical rigor.¹¹ Since we are looking to train an event categorizer, let us look at the

¹⁰Since the total number of events is itself a random variable with different distributions under $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$, the mere occurrence/existence of an event in data also carries some information relevant to distinguishing between $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$, in addition to $p(\mathbf{e})$. But in the context of event categorization, this is a moot point.

¹¹The following argument can be presented more rigorously by considering a probabilistic event categorizer that returns the probabilities of assigning an event to different categories, and using functional derivatives to capture *small* changes to the categorizer. This, however, is unnecessary for our purposes here.

effect of modifying a given event categorizer ‘slightly’. Let us modify an event categorizer so that a given event $\mathbf{e}$ moves from category c_1 to category c_2 . This changes $D_{\text{Neym}\chi^2}$ by changing $s_{c_1}^2(\mathbf{x})/n_{c_1}(\mathbf{x})$ and $s_{c_2}^2(\mathbf{x})/n_{c_2}(\mathbf{x})$ for $\mathbf{x} = \mathbf{x}(\mathbf{e})$.

$$\Delta \left(\frac{s_c^2(\mathbf{x})}{n_c(\mathbf{x})} \right) = -\frac{s_c^2(\mathbf{x})}{n_c^2(\mathbf{x})} \Delta n_c(\mathbf{x}) + 2 \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \Delta s_c(\mathbf{x}) \quad (2.3)$$

$$= \Delta n_c(\mathbf{x}) \left[-p_c^2(\mathbf{x}) + 2p_c(\mathbf{x})p(\mathbf{e}) \right] \quad (2.4)$$

Here $\Delta n_c(\mathbf{x})$ and $\Delta s_c(\mathbf{x})$ represent small changes to $n_c(\mathbf{x})$ and $s_c(\mathbf{x})$ caused by the change in the event categorizer, with $\Delta s_c(\mathbf{x}) = p(\mathbf{e})\Delta n_c(\mathbf{x})$. Note that for our modification of moving event $\mathbf{e}$ from c_1 to c_2 , $\Delta n_{c_2} = -\Delta n_{c_1}$. The first term captures the loss (gain) from increasing (decreasing) n_c , and the second term captures the gain (loss) from increasing (decreasing) s_c . Armed with this intuition, let us rewrite $D_{\text{Neym}\chi^2}$ as follows:

$$D_{\text{Neym}\chi^2} = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \, n_c(\mathbf{x}) \frac{s_c^2(\mathbf{x})}{n_c^2(\mathbf{x})} \quad (2.5)$$

$$= \sum_{c=1}^C \int_{\Omega_{\mathbf{e}}} d\mathbf{e} \, n(\mathbf{e}) \, \delta(c, \eta(\mathbf{e})) \frac{s_c^2(\mathbf{x})}{n_c^2(\mathbf{x})} \quad (2.6)$$

$$= \int_{\Omega_{\mathbf{e}}} d\mathbf{e} \, n(\mathbf{e}) \sum_{c=1}^C \delta(c, \eta(\mathbf{e})) \frac{s_c^2(\mathbf{x})}{n_c^2(\mathbf{x})} \quad (2.7)$$

$$= \int_{\Omega_{\mathbf{e}}} d\mathbf{e} \, n(\mathbf{e}) \sum_{c=1}^C \delta(c, \eta(\mathbf{e})) \left[-\frac{s_c^2(\mathbf{x})}{n_c^2(\mathbf{x})} + 2 \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \frac{s(\mathbf{e})}{n(\mathbf{e})} \right] \quad (2.8)$$

$$= \int_{\Omega_{\mathbf{e}}} d\mathbf{e} \, n(\mathbf{e}) \sum_{c=1}^C \delta(c, \eta(\mathbf{e})) \left[-p_c^2(\mathbf{x}) + 2p_c(\mathbf{x})p(\mathbf{e}) \right] \quad (2.9)$$

$$= N \times \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \eta(\mathbf{e})) \left[-p_c^2(\mathbf{x}) + 2p_c(\mathbf{x})p(\mathbf{e}) \right] \right], \quad (2.10)$$

where δ is the Kronecker delta, and $\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} [\dots]$ represents the expected value computed over events $\mathbf{e}$ sampled from $\mathbb{H}_{b+s}$. As a reminder, in eqs. (2.6)–(2.10), $\mathbf{x}$ represents $\mathbf{x}(\mathbf{e})$ of the event $\mathbf{e}$ being integrated over. Note the commonality between the right-hand sides of eqs. (2.4) and (2.10).

We will first justify each step of the derivation, and then unpack the power of the last equation for our categorizer training purpose. Eq. (2.6) follows from the fact that

$$n_c(\mathbf{x}') = \int_{\Omega_{\mathbf{e}}} d\mathbf{e} \, n(\mathbf{e}) \, \delta(c, \eta(\mathbf{e})) \, \delta_{\mathbf{x}}(\mathbf{x}' - \mathbf{x}(\mathbf{e})), \quad (2.11)$$

where δ is the Kronecker delta function, and $\delta_{\mathbf{x}}$ satisfies¹²

$$\int_{\Omega_{\mathbf{x}}} d\mathbf{x}' \, \delta_{\mathbf{x}}(\mathbf{x}' - \mathbf{x}) \, f(\mathbf{x}') = f(\mathbf{x}) \quad (2.12)$$

¹² $\delta_{\mathbf{x}}$ can be thought of as a product of Kronecker deltas for the discrete components of $\mathbf{x}$, and Dirac delta functions for the real valued components.

Eq. (2.8) follows from the fact that $s_c(\mathbf{x})/n_c(\mathbf{x})$ is the expected value of $s(e)/n(e)$ over events sampled from $\mathbb{H}_{b+s}$, conditional on their c and $\mathbf{x}$ values. More concretely, eq. (2.8) can be shown to follow from

$$\frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} = \frac{1}{n_c(\mathbf{x})} \int_{\Omega_e} d\mathbf{e}' n(\mathbf{e}') \delta(c, \eta(\mathbf{e}')) \delta_{\mathbf{x}}(\mathbf{x} - \mathbf{x}(\mathbf{e}')) \frac{s(\mathbf{e}')}{n(\mathbf{e}')} . \quad (2.13)$$

Eq. (2.9) follows from the definitions of $p_c(\mathbf{x})$ and $p(e)$, and eq. (2.10) follows from the definition of expected value of a random variable over events sampled from a certain distribution.

Based on the form of eq. (2.10), let us define a function $\Lambda_{\text{Neym}\chi^2} : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ as

$$\Lambda_{\text{Neym}\chi^2}(U, V) \equiv -U^2 + 2UV . \quad (2.14)$$

$\Lambda_{\text{Neym}\chi^2}(U, V)$ represents the eventwise contribution to the statistical distance $D_{\text{Neym}\chi^2}$ of adding an event e with $p(e) = V$ to a category c with $p_c(\mathbf{x}(e)) = U$. This lets us write eq. (2.10) as

$$\frac{D_{\text{Neym}\chi^2}}{N} = \text{MEAN}_{e \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \eta(e)) \Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(e)) \right] . \quad (2.15)$$

To understand the meaning of the phrase “eventwise contribution to the statistical distance”, let us see how eq. (2.15) can be used to estimate $D_{\text{Neym}\chi^2} / N$ for a given event categorizer η .

Estimating $D_{\text{Neym}\chi^2} / N$ for a *given event categorizer* η :

1. Use the event categorizer η to categorize events sampled from $\mathbb{H}_{b+s}$.
2. Estimate $p_c(\mathbf{x})$ for all c and $\mathbf{x}$ using the categorized events. This can be done either as the average of $p(e)$ conditional on $(c, \mathbf{x})$, or directly from the label carried by each training event indicating whether it is a signal or background event.
3. Next for each event e , estimate its contribution to $D_{\text{Neym}\chi^2}$ as $\Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(e))$, where c is the category it belongs to *under the given event categorizer* η . If event rejection is allowed, the contribution from rejected events is just 0. The average value of this contribution over training events is an estimate for $D_{\text{Neym}\chi^2} / N$.¹³

In this way, the $\Lambda_{\text{Neym}\chi^2}$ term represents the eventwise contribution to $D_{\text{Neym}\chi^2}$. While this is admittedly a roundabout way of estimating $D_{\text{Neym}\chi^2}$ from a training sample, it leads to a method of training an optimal event categorizer iteratively, as we will see next.

2.4 Iterative maximization of $D_{\text{Neym}\chi^2}$

The idea behind writing the statistical distance, $D_{\text{Neym}\chi^2}$, in terms of an eventwise contribution, $\Lambda_{\text{Neym}\chi^2}$, is that from event categorizer η , we can get a “better” event categorizer $\bar{\eta}$ by requiring that $\bar{\eta}$ places each event in the category that maximizes $\Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(e))$ (with $p_c(\mathbf{x})$ corresponding to the categorizer η).

This leads to the following iterative prescription to train an event categorizer:

¹³Note that N is the expected number of events in the real experiment under $\mathbb{H}_{b+s}$, not the number of training events.

Figure 3. A cartoon depiction of $\mathbf{x}$ dependent boundaries on $p(\mathbf{e})$ between different categories.

1. Let the current categorizer be η . Categorize events sampled from $\mathbb{H}_{b+s}$ as per η .
2. Estimate $p_c(\mathbf{x})$ for all c and $\mathbf{x}$ (for the current η) using the categorized events.
3. Construct a new categorizer $\bar{\eta}$ which places each event $\mathbf{e}$ in the category c with the highest $\Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(\mathbf{e}))$. If event rejection is allowed, then an event will be rejected by $\bar{\eta}$, iff $\Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(\mathbf{e})) < 0$ for all categories c .
4. Set $\eta = \bar{\eta}$. Repeat from step 1. Stop using some appropriate termination condition.

We will postpone the detailed description of this training procedure to section 3, which is intended to serve as a usage guide independent from the build-up we are doing now.

But first, to gain some intuition, let us take a closer look at what the iterative prescription described above does operationally. From a given categorizer η with corresponding signal-fraction functions $p_c(\mathbf{x})$, we build a new categorizer $\bar{\eta}$ as¹⁴

$$\bar{\eta}(\mathbf{e}) = \arg \max_c \left[\Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(\mathbf{e})) \right] , \quad (2.16)$$

where $\Lambda_{\text{Neym}\chi^2}(U, V) = -U^2 + 2UV$. Note that $\Lambda_{\text{Neym}\chi^2}(U, V)$ has a global maximum in the first argument at $U = V$, and monotonically decreases as U moves away from V . It is also symmetric¹⁵ in U about this maximum. Combining this observation with eq. (2.16), we can see that $\bar{\eta}$ places each event in the category with $p_c(\mathbf{x})$ closest to $p(\mathbf{e})$. We can think of $\bar{\eta}$ as grouping events by their $p(\mathbf{e})$ values with $\mathbf{x}$ dependent boundaries between different categories — figure 3 shows a cartoon of what such a categorization could look like. This makes sense considering that the purpose of categorization is to reduce the information loss arising from mixing up regions of phase-space with different $p(\mathbf{e})$ values when reducing the full event $\mathbf{e}$ to the event variable $\mathbf{x}$.

¹⁴To keep the presentation simple, we are ignoring the case where event rejection is allowed. $\bar{\eta}(\mathbf{e})$ will be *rejected* or *None* if every category leads to a negative contribution.

¹⁵Maximum in the first argument at $U = V$ and monotonic decrease as U moves away from this maximum will hold for the eventwise contributions for other statistical distances as well, but not the symmetry.

Recall that $p_c(\mathbf{x})$ is the mean p in category c conditional on $\mathbf{x}$. This process of iteratively placing events in the category with the closest mean, and then recomputing the mean for each category could be reminiscent of Lloyd’s k-means clustering algorithm [45, 46]. This connection can be made clearer by noting from eq. (2.14) that

$$\Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(\mathbf{e})) = -[p(\mathbf{e}) - p_c(\mathbf{x})]^2 + p(\mathbf{e})^2 \quad . \quad (2.17)$$

The first term is congruous to $[X - \langle X \rangle_c]^2$, the objective function k-means clustering seeks to minimize the average of, where $\langle X \rangle_c$ is the sample mean within cluster c which contains X , and the last term is independent of the event categorization criteria. More explicitly, the objective function to be minimized in k-means clustering can be written as

$$\sum_X \sum_{\text{cluster } c} \delta \left(c, \begin{array}{c} \text{cluster containing} \\ X \end{array} \right) [X - \langle X \rangle_c]^2 \quad , \quad (2.18)$$

which bears close resemblance to the right-hand side of eq. (2.15). In fact, for the case of $D_{\text{Neym}\chi^2}$, our iterative training procedure can be thought of as the k-means clustering algorithm with events within each $\mathbf{x}$ -“bin” being clustered based on their $p(\mathbf{e})$ values. But instead of doing this bin-by-bin in $\mathbf{x}$, we can take an unbinned approach to the estimation of $p_c(\mathbf{x})$. We will elaborate more on this in section 3.

Note that we have not yet proved that the categorizer will “improve” during each iterative step or that it will converge on an optimal categorizer in a finite number of steps — at this point the algorithm is held together by the heuristic argument presented in eq. (2.4) and the connection to k-means clustering. In appendix C, we present a proof of the correctness of the iterative training algorithm (for $D_{\text{Neym}\chi^2}$, as well as the other statistical distances). The proof will be similar to that of the generic k-means clustering algorithm, and will rely on the fact that $\Lambda_{\text{Neym}\chi^2}(p_c(\mathbf{x}), p(\mathbf{e}))$ and its counterparts for the other distances are linear in $p(\mathbf{e})$ and have a global maximum in the first argument at $p_c(\mathbf{x}) = p(\mathbf{e})$.

2.5 Generalizing to other statistical distances

We can repeat the process in section 2.3 for the other statistical distances in eq. (2.1) and generate eventwise contributions to the other statistical distances as well. We will skip the derivation and only present the relevant results here. First let us define $\tilde{D}_{\text{stat}}$ as

$$\tilde{D}_{\text{Neym}\chi^2} \equiv \frac{D_{\text{Neym}\chi^2}}{N} \quad , \quad (2.19a)$$

$$\tilde{D}_{\text{Pear}\chi^2} \equiv \frac{D_{\text{Pear}\chi^2}}{N} \quad , \quad (2.19b)$$

$$\tilde{D}_{\text{KL}} \equiv 2 \frac{D_{\text{KL}}}{N} \quad , \quad (2.19c)$$

$$\tilde{D}_{\text{revKL}} \equiv 2 \frac{D_{\text{revKL}}}{N} \quad , \quad (2.19d)$$

$$\tilde{D}_{\text{J}} \equiv \frac{D_{\text{J}}}{N} \quad , \quad (2.19e)$$

$$\tilde{D}_B \equiv 8 \frac{D_B}{N} . \quad (2.19f)$$

The numerical factors of 2 and 8 are introduced to make all the $\tilde{D}_{\text{stat}}$ -s converge to each other in the limit $s_c(\mathbf{x})/n_c(\mathbf{x}) \rightarrow 0$ for all c and almost all $\mathbf{x}$. For all $\tilde{D}_{\text{stat}}$ -s, we have

$$\tilde{D}_{\text{stat}} = \text{MEAN}_{e \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \eta(e)) \Lambda_{\text{stat}}(p_c(\mathbf{x}), p(e)) \right] , \quad (2.20)$$

where the functions $\Lambda_{\text{stat}} : [0, 1) \times [0, 1] \rightarrow \mathbb{R}$ are given by¹⁶

$$\Lambda_{\text{Neym}\chi^2}(U, V) = -U^2 + 2UV , \quad (2.21a)$$

$$\Lambda_{\text{Pear}\chi^2}(U, V) = -\frac{U^2}{(1-U)^2} + \frac{(2-U)UV}{(1-U)^2} , \quad (2.21b)$$

$$\Lambda_{\text{KL}}(U, V) = -2 \left[\ln(1-U) + \frac{U}{1-U} \right] + \frac{2UV}{1-U} , \quad (2.21c)$$

$$\Lambda_{\text{revKL}}(U, V) = 2[\ln(1-U) + U] - 2\ln(1-U)V , \quad (2.21d)$$

$$\Lambda_J(U, V) = -\frac{U^2}{1-U} + \left[\frac{U}{1-U} - \ln(1-U) \right] V , \quad (2.21e)$$

$$\Lambda_B(U, V) = 4 \left[2 - \frac{2-U}{\sqrt{1-U}} \right] + 4 \left[\frac{1}{\sqrt{1-U}} - 1 \right] V . \quad (2.21f)$$

With these choices, the validity of eq. (2.20) can be verified (after some manipulation) by noting that because $\Lambda_{\text{stat}}(p_c(\mathbf{x}), p(e))$ is linear in $p(e)$, it can be replaced by $\Lambda_{\text{stat}}(p_c(\mathbf{x}), p_c(\mathbf{x}))$ in eq. (2.20).

For each Λ_{stat} it can be shown that for all $V \in [0, 1]$, $\Lambda_{\text{stat}}(U, V)$ is monotonically increasing in U for $0 \leq U \leq V$ and monotonically decreasing in U for $V \leq U < 1$, with global maximum in U attained at $U = V$. This means that in each iterative step, the “next” categorizer will place event e in one of the two categories with “current” $p_c(\mathbf{x})$ closest to $p(e)$ on either side, or reject the event if that option is allowed.

2.6 The rectified “misalignment”

In the introductory section 1.1, the *source* of misalignment between the physics goals and typical ML-based event selectors was attributed to not ensuring that categorization is based only on information complementary to the event variable $\mathbf{x}$. Now we are in a position to point out *what* that misalignment is. In typical ML-based selector training, all signal events are treated on the same footing and all background events are treated on the same footing, in the sense that the reward/penalty for including/excluding an event depends only on whether the event is from the signal or background sample. On the other hand, our eventwise contribution functions $\Lambda_{\text{stat}}(p_c(\mathbf{x}), p(e))$ suggest that the reward/penalty should also depend on the $\mathbf{x}$ value of the event. For the rest of this subsection, let us restrict our attention to event selectors only, for simplicity.

¹⁶ Λ_{KL} with $U = V = s(e)/n(e)$ is related to the per-event significance indicator ξ introduced in [47].

Figure 4. The penalty (left panel) and the reward (right panel) associated with selecting a background and a signal event, respectively, into an $\mathbf{x}$ -bin of purity p_c . The y -axes of the plots are in log-scale.

Negative $\Lambda_{\text{stat}}(p_c, 0)$ represents the penalty for selecting a background event in an $\mathbf{x}$ -“bin” of signal purity p_c . As can be seen from the left panel of figure 4, this penalty increases with p_c for each of our distance measures. Similarly, $\Lambda_{\text{stat}}(p_c, 1)$ represents the reward for selecting a signal event in an $\mathbf{x}$ -“bin” of signal purity p_c . As can be seen from the right panel of figure 4, this reward also increases with p_c .

An event $\mathbf{e}$ will warrant selection only if the following net-reward corresponding to $p(\mathbf{e})$ is non-negative:

$$\text{net-reward} = (1 - p(\mathbf{e}))\Lambda_{\text{stat}}(p_c(\mathbf{x}), 0) + p(\mathbf{e})\Lambda_{\text{stat}}(p_c(\mathbf{x}), 1) \quad (2.22)$$

This net-reward is precisely $\Lambda_{\text{stat}}(p_c(\mathbf{x}), p(\mathbf{e}))$. Let $p_{\text{stat}}^{\text{sel_thresh}}(p_c)$ be the selection threshold on $p(\mathbf{e})$ above which an event warrants being selected into a bin of purity p_c . This threshold is defined by the equation

$$\Lambda_{\text{stat}}[p_c, p_{\text{stat}}^{\text{sel_thresh}}(p_c)] = 0 \quad . \quad (2.23)$$

Figure 5 shows that this threshold is an increasing function of p_c for each of our statistical distances, implying that the threshold is stronger in $\mathbf{x}$ -“bins” of higher purity. This is how ThickBrick naturally leads to decorrelation of the event selection criteria from the event variable $\mathbf{x}$.

Failure to incorporate these $\mathbf{x}$ dependences, or more precisely the $p_c(\mathbf{x})$ dependences, is the misalignment present in typical ML-based selector training that our ThickBrick prescription rectifies. The fact that $p_c(\mathbf{x})$ itself depends on the event selector is the reason our training prescription changes the per-event reward function being maximized at each iterative step.

3 Prescription for training an optimal event categorizer

In this section we provide a description of our training prescription. Operationally, eqs. (2.20) and (2.21), along with the definitions of $p(\mathbf{e})$ and $p_c(\mathbf{x})$ in table 1, capture

Figure 5. The selection threshold $p_{\text{stat}}^{\text{sel_thresh}}(p_c)$ for the different statistical distances.

almost all the information we need to carry from the previous section (groundwork) into this one.

Let us begin with a sample of simulated events containing both signal and background, split into training and validation sets. For every simulated event i :

- Let $\mathbf{e}_i$ represent all the observable information from the event.
- Let w_i be the non-negative weight associated with the event.¹⁷
- Let y_i be the label specifying the subsample the event is from: 1 for signal, 0 for background. In HEP applications, this labeling is trivially possible because we know which model was used to simulate a particular event.
- Let $\mathbf{x}_i = \mathbf{x}(\mathbf{e}_i)$ be the value of the event variable $\mathbf{x}$ whose distribution is to be analyzed to search for the presence of a signal.

Let $\tilde{S}/\tilde{B}$ be the overall signal to background ratio in the *training* sample. We will *not* assume that this matches the overall signal to background ratio S/B expected in an experiment under the hypothesis $\mathbb{H}_{b+s}$. This is to allow for the fact that this ratio could be a priori unknown in the search (if the signal cross-section is an unspecified parameter in the theory). Even in cases when the ratio S/B is known, working with a highly imbalanced sample with very different values for $\tilde{S}$ and $\tilde{B}$ can make learning $p(\mathbf{e})$ using ML techniques difficult. So it is often preferable to use training samples with $\tilde{S}/\tilde{B}$ close to 1.

3.1 Preprocessing: learning $p(\mathbf{e})$

From the training data we can estimate $p(\mathbf{e})$ as the expected value of y_i , weighted by w_i , conditional on $\mathbf{e}$. If the dimensionality of $\mathbf{e}$ is large, this can be done with machine learning techniques [14, 17], or the Matrix Element Method [48, 49]. The accuracy of the estimate

¹⁷We accommodate the possibility of weighted events in order to allow for easy scans of BSM parameter spaces by simple reweighting of training events using the matrix element [2, 48, 49].

will depend on several factors including, but not limited to, the amount of computational resources and training data available.

We can address the limitation posed by inaccuracies in the estimation of $p(\mathbf{e})$ by treating the output of the classifying algorithm (BDT or neural network for example) not as $p(\mathbf{e})$ but simply as a property of the event to base our categorization on. Let $O_{\text{ML}}(\mathbf{e})$ be the machine learning classifier output. We can estimate $p(\mathbf{e})$ as the probability that an event sampled from $\mathbb{H}_{b+s}$ will be a signal event conditional on *both* $O_{\text{ML}}(\mathbf{e})$ and $\mathbf{x}(\mathbf{e})$, as opposed to conditional on $\mathbf{e}$. Presumably the dimensionality of $\mathbf{x}(\mathbf{e})$ is small enough to perform this estimation with sufficient accuracy. This can be done either using traditional regression methods (parameterized or unparameterized), or with another machine learning layer (with O_{ML} and $\mathbf{x}$ as inputs). Henceforth, $p(\mathbf{e}_i)$ and p_i will both refer to this estimate of $p(\mathbf{e}_i)$.¹⁸

Out of the six statistical distance measures, only $\tilde{D}_{\text{Pear}\chi^2}$ maximization leads to a categorizer that is independent of the value of the overall signal to background ratio S/B under $\mathbb{H}_{b+s}$. Training using the other five measures is sensitive to S/B ; if training is to be done using these measures, w_i and p_i both need to be appropriately scaled before proceeding further. To scale the signal cross-section by a factor of λ (i.e., if $S/B = \lambda\tilde{S}/\tilde{B}$), we need to apply the following transformations¹⁹

$$w_i \rightarrow \lambda w_i \quad \text{if } y_i = 1, \text{ i.e., for signal events only}^{20} \quad (3.1a)$$

$$p_i \rightarrow \frac{\lambda p_i}{1 + (\lambda - 1)p_i} \quad \text{for all events } i \quad (3.1b)$$

In searches where there is no a priori expectation for the size of the signal cross-section,²¹ one option is to use $\tilde{D}_{\text{Pear}\chi^2}$ as the distance measure or choose a very small value of λ (which is equivalent to using $\tilde{D}_{\text{Pear}\chi^2}$). Another option is to choose a value of λ that puts the signal cross-section either near the low end of discoverability or near the potential upper-limit on the signal cross-section (low end of rejectability) expected from the available experimental data.

At the end of this preprocessing, each training event should have w , y , $\mathbf{x}$, and p associated with it (possibly scaled). In addition, we should have a way of computing $\mathbf{x}$ and p for validation MC events as well as real events from the experiment (by first running the event $\mathbf{e}$ through the machine learning classifier to get $O_{\text{ML}}(\mathbf{e})$ and using that to get p , performing the scaling operations in eq. (3.1) if applicable).

It is important that the following equation approximately holds true for validation MC events sampled from $\mathbb{H}_{b+s}$:

$$p \approx \frac{\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[wy \mid \mathbf{x}, p \right]}{\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[w \mid \mathbf{x}, p \right]} \quad (3.2)$$

¹⁸For convenience, we are not introducing new notation, like $\hat{p}$, to refer to the *estimate* of p .

¹⁹It is possible to postpone this scaling and handle it during training, which is what is done internally in our python implementation of ThickBrick [38].

²⁰Alternatively, the transformation $w_i \rightarrow (1 + (\lambda - 1)p_i)w_i$ can be applied to all events.

²¹SUSY and UED searches would be examples where there exists an a priori expectation for the signal cross-section, which is governed by the known gauge couplings and BSM partner spins.

Eq. (3.2) just means that p approximately equals the probability of an event sampled from $\mathbb{H}_{b+s}$ being a signal event, conditional on $(\mathbf{x}, p)$, with the inexactness resulting from training limitations.

3.2 Choices to be made

1. *Choose a distance measure to train with.* As mentioned earlier, ThickBrick involves choosing a distance measure between the ‘background only’ and ‘background+signal’ hypotheses to be maximized by the event categorizer. The possible options are $\tilde{D}_{\text{Neym}\chi^2}$, $\tilde{D}_{\text{Pear}\chi^2}$, $\tilde{D}_{\text{KL}}$, $\tilde{D}_{\text{revKL}}$, $\tilde{D}_{\text{J}}$, and $\tilde{D}_{\text{B}}$.

In actual experiments, previously undetected signals (if present) are expected to be faint. In this limit, all the statistical distances converge to one another up to constant factors as-per (2.2). As a result, the choice of distance measure will be of little consequence in most situations. However, here we will comment on the differences between the distance measures (when the signal being searched is not weak) for academic purposes.

As mentioned above in section 3.1, maximizing $\tilde{D}_{\text{Pear}\chi^2}$ leads to a categorizer whose optimality is independent of the overall signal to background ratio. This is because scaling the signal number density by *an overall factor* ($s(\mathbf{e}) \rightarrow \lambda s(\mathbf{e})$) only changes $D_{\text{Pear}\chi^2}$ by a multiplicative factor ($D_{\text{Pear}\chi^2} \rightarrow \lambda^2 D_{\text{Pear}\chi^2}$). This makes $\tilde{D}_{\text{Pear}\chi^2}$ a good option when there is a priori no expectation for the potential signal cross-section.

We show in appendix B that training using $\tilde{D}_{\text{revKL}}$ (with training signal cross-section near the low end of discoverability) is theoretically optimal for signal discovery. Similarly, training using $\tilde{D}_{\text{KL}}$ (with training signal cross-section near the expected upper-limit on it) is theoretically optimal for upper-limit setting.

In practice, the main difference between the different distance measures lies in the ($\mathbf{x}$ dependent) boundaries on the values of p between different categories — they have different thresholds to move events from a low purity category to a high purity category. Let us call a distance measure that leads to higher thresholds as ‘stricter’. We show in appendix D that the distance measures arranged in the increasing order of strictness are as follows (left to right, least strict to most strict):

$$\tilde{D}_{\text{Neym}\chi^2} \prec \tilde{D}_{\text{revKL}} \prec \tilde{D}_{\text{B}} \prec \tilde{D}_{\text{J}} \prec \tilde{D}_{\text{KL}} \prec \tilde{D}_{\text{Pear}\chi^2} \quad (3.3)$$

Here, the notation $\tilde{D}_{\text{Neym}\chi^2} \prec \tilde{D}_{\text{revKL}}$ just means that $\tilde{D}_{\text{Neym}\chi^2}$ precedes $\tilde{D}_{\text{revKL}}$ in order of strictness (as described above). Note that the hierarchy (3.3) is reflected in the ordering of the selection thresholds in figure 5.

One might want to try out different distance measures, and pick one based on practical considerations like having sufficient number of events in each category, shape features of the distribution of $\mathbf{x}$ post-categorization, etc. — it is possible that background MC validation (post-categorization) is easier when using certain distance measures than others.

Henceforth, where applicable, $\tilde{D}_{\text{stat}}$ will refer to the chosen distance measure.

2. Choose the number of categories (`int C ≥ 1`) and whether or not to allow event rejection (`bool allow_rejection`).

To use our prescription to train an event selector (as opposed to a categorizer), set the number of categories `C` to 1, and set `allow_rejection` to `True`.

Theoretically, in the asymptotic (large number of events) limit, having more categories will lead to higher sensitivity. In fact, in the infinite categories limit, the sensitivity of the search will approach the sensitivity of a likelihood-ratio test performed in the full phase-space $\mathbf{e}$, which is the theoretical upper-limit of the sensitivity of the search. Also, for a given number of categories, allowing event rejection will lead to higher sensitivity than disallowing it. In practice, however, increasing the number of categories will lead to data getting diluted out across categories, which can lead to problems.

Fortunately, though, continually increasing the number of categories typically leads to diminishing returns.²² In section 3.4 we will show how to estimate the value of $\tilde{D}_{\text{stat}}$ (post training) in the $C \rightarrow \infty$ limit. This can be used to assess the utility of increasing the number of categories and weigh it against the corresponding disadvantages, in order to settle on a good value of C .

3.3 Iterative training prescription

Algorithm 1 shows the iterative event categorizer training prescription introduced in section 2.4. We have already motivated the algorithm in section 2, and the proof of correctness has been relegated to appendix C. Here, we will elaborate on different parts of the algorithm from an implementation point of view, as well as provide diagnostic notes from a practical usage point of view.

Lines 1 and 2 refer to the preprocessing covered in section 3.1 and the choices covered in section 3.2.

Initialization (line 3): the categorizer η will categorize events based on their $\mathbf{x}$ and p values. It can be initialized to apply arbitrary cutoffs/thresholds on the p values. The initial cutoffs/thresholds could possibly be $\mathbf{x}$ dependent, though this is not necessary. η can also be initialized to a randomized categorizer or to a constant function.

Note that η being a function of $\mathbf{x}$ and p is consistent with our original notation of η being a function of event $\mathbf{e}$, since $\mathbf{x}(\mathbf{e})$ and $p(\mathbf{e})$ are part of the event description after the preprocessing step. Also note that in our description here, η is expected to return an integer between 1 and C or, alternatively, “rejected” if `allow_rejection` is set to `True`.

Lines 5–21 contain the code block for one iterative training step.

²²The diminishing returns can be attributed to the differentiable behavior of $\Lambda_{\text{stat}}(U, V)$ in parameter U at/near the maximum $U = V$. As a result, the returns from creating more categories to reduce the distance between $p(\mathbf{e})$ and the closest $p_c(\mathbf{x})$ are diminishing.

	Result:	Trained categorizer η
1		Preprocess training events (get w , y , $\mathbf{x}$, and p for all events)
2		Choose $\tilde{D}_{\text{stat}}$, C , <code>allow_rejection</code>
3	initialize:	categorizer η , a function of $\mathbf{x}$ and p
4	repeat	
		*/
5		Categorize events in the training set based on their $\mathbf{x}$ and p values, as per η
6		Estimate the functions $p_c(\mathbf{x})$ for $c \in \{1, \dots, C\}$ (for current η)
7	define	<code>Next_categorizer($\mathbf{x}$, p):</code>
8		<code>$p_{\text{c_list}} = \text{sorted}([p_c(\mathbf{x}) \text{ for } c \in \{1, \dots, C\}])$</code>
9		<code>$\text{max_}\Lambda = \max_{i \in \{1, \dots, C\}} [\Lambda_{\text{stat}}[p_{\text{c_list}}[i], p]]$</code>
10		<code>$\text{best_c} = \arg \max_{i \in \{1, \dots, C\}} [\Lambda_{\text{stat}}[p_{\text{c_list}}[i], p]]$</code>
11		if (<code>$\text{max_}\Lambda \geq 0$</code>) then
12		<code>return best_c</code>
13		else
14		if (<code>$\text{allow_rejection} == \text{True}$</code>) then
15		<code>return rejected</code>
16		else
17		<code>return best_c</code>
18		end
19		end
20	end definition	
21	update:	$\eta = \text{Next_categorizer}$
22	until	appropriate termination condition based on training convergence

Algorithm 1. ThickBrick’s iterative prescription to train η .

Estimation of $p_c(\mathbf{x})$ (lines 5, 6):²³ to estimate the signal-fraction functions $p_c(\mathbf{x})$ for a given categorizer η , training events are first categorized based on η . Then, from training events in a given category c , $p_c(\mathbf{x})$ can be estimated as the expected value of either p or y (weighted by w) conditional on $\mathbf{x}$. This regression problem can be solved using a plethora of techniques — binned and unbinned, parametric and non-parametric (like kernel regression).

The estimation of $p_c(\mathbf{x})$ for a given categorizer comes with the typical challenges of machine learning, including the bias-variance trade-off. One might want to employ cross-validation techniques (for example, in picking the bandwidth for kernel regression) to get the best results. If the estimation of $p_c(\mathbf{x})$ is affected in some regions of $\mathbf{x}$ by a low density of training events, it might be helpful to increase the sampling density in those regions

²³Note that the estimation of $p_c(\mathbf{x})$ could be split between line 6 and line 8, with line 6 representing any required preprocessing and line 8 representing the actual estimation for particular values of $\mathbf{x}$ and c . This is typically the case in kernel regression.

(and appropriately decrease the event weights). Of course, with a reasonably large number of training events, and a reasonably small dimensionality of $\mathbf{x}$ and number of categories, these measures will be rendered inessential.

Defining the next categorizer (lines 7–20): the next categorizer is constructed so as to reject events if $\Lambda_{\text{stat}}(p_c(\mathbf{x}), p) < 0$ for all categories c (and if event rejection is allowed). Otherwise, the next categorizer groups events at a given value of $\mathbf{x}$ into different categories, based on their p values.

The sorting in line 8 ensures that at each value of $\mathbf{x}$, events in a *higher* category will have higher value of p . In other words, line 8 ensures that category boundaries (on values of p) do not interweave or cross each other when $\mathbf{x}$ is varied. Note that, our distance measures $\tilde{D}_{\text{stat}}$ are all invariant under an $\mathbf{x}$ dependent cyclic permutation of categories. But having the categories ordered by their p_c at each $\mathbf{x}$ makes the categorization more meaningful. This also helps with the estimation of the signal-fraction functions $p_c(\mathbf{x})$ by reducing their sensitivity to changes in $\mathbf{x}$.

Note that during each iteration, the next categorizer depends implicitly on the estimated functions, $p_c(\mathbf{x})$, for the current categorizer. This means that the estimated functions $p_c(\mathbf{x})$ will need to be packaged into the final product of the training, which is the trained categorizer. The details of this packaging will depend on the particular regression technique employed.

Termination condition (line 21): Lloyd’s k-means clustering algorithm converges and terminates when the category assignments do not change in a training step. This convergence is guaranteed within a finite number of steps, and in practice the fraction of events whose category assignments change during a step drops significantly over only a few tens of iterations.

The guarantee of convergence within a finite number of steps extends to our algorithm as well (as argued in appendix C) if the estimation of $p_c(\mathbf{x})$ is performed by binning events into different (fixed) $\mathbf{x}$ -bins and finding the expected value of p or y in each bin. But if one takes an unbinned approach to the estimation of $p_c(\mathbf{x})$, it is possible to run into cycles near the end of the training process due to artifacts in the regression technique. This can be handled by terminating the iterative training when the fraction of events reassigned in some step drops below a certain threshold.

One can also terminate the algorithm after a predetermined number of steps — the convergence of the algorithm will typically be fast enough (for small values of C) to warrant this. Alternative to these, we can also employ a stopping condition based on the saturation of $\tilde{D}_{\text{stat}}$, as explained below.

3.4 Performance evaluation

For any given categorizer η (intermediate or final), $\tilde{D}_{\text{stat}}$ can be estimated from eq. (2.20), after first estimating the signal-fraction functions $p_c(\mathbf{x})$. It can also be estimated as

$$\tilde{D}_{\text{stat}} = \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \eta(\mathbf{e})) \Lambda_{\text{stat}}(p_c(\mathbf{x}), p_c(\mathbf{x})) \right] . \quad (3.4)$$

This can serve multiple purposes. First, we can use it to devise a termination condition — terminate the training when the fractional or absolute increase in the (estimated) value of $\tilde{D}_{\text{stat}}$ drops below a certain threshold.²⁴

Secondly, the values of $\tilde{D}_{\text{stat}}$ for categorizers trained with different values of $\mathbb{C}$ (estimated from validation datasets) can be used to choose between different values of $\mathbb{C}$. Furthermore, in the $\mathbb{C} \rightarrow \infty$ limit, the value of $\tilde{D}_{\text{stat}}$ (for the optimal categorizer) can be estimated as

$$\lim_{\mathbb{C} \rightarrow \infty} (\text{optimal } \tilde{D}_{\text{stat}}) = \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} [\Lambda_{\text{stat}}(p(\mathbf{e}), p(\mathbf{e}))] \quad . \quad (3.5)$$

This is also theoretical upper-limit on $\tilde{D}_{\text{stat}}$ achievable by analyzing the distribution of events in the full-phase space $\mathbf{e}$. As mentioned in section 3.2 this can be used to assess the improvement one stands to gain by continually increasing $\mathbb{C}$. Eq. (3.5) follows from eq. (2.20) by noting that the optimal categorizer in the infinite categories limit will keep events with different values of p from mixing.

Note that comparisons between categorizers based on $\tilde{D}_{\text{stat}}$ will only be meaningful if a common performance measure is employed, regardless of the performance measure the individual categorizers were trained to maximize. In fact, it can be shown that for any given $\mathbb{H}_b$, $\mathbb{H}_{b+s}$, and categorizer η , the different $\tilde{D}_{\text{stat}}$ values will satisfy

$$\tilde{D}_{\text{Neym}\chi^2} \leq \tilde{D}_{\text{revKL}} \leq \tilde{D}_{\text{B}} \leq \tilde{D}_{\text{J}} \leq \tilde{D}_{\text{KL}} \leq \tilde{D}_{\text{Pear}\chi^2} \quad , \quad (3.6)$$

which also happens to be the same as the order of strictness in eq. (3.3).

3.5 Miscellaneous notes

- Note that the iterative training procedure will converge on a *locally optimal* categorizer. To ensure that the training does not lead to a categorizer much worse than the globally optimal one, one may have to repeat the training procedure with different initial categorizers.
- It is possible that one is not interested in $\mathbf{x}$ dependent cutoffs and thresholds on p , and simply wants to maximize the overall $\sum_c S_c^2/N_c$ (or the equivalent for a different statistical distance), using flat boundaries between categories. In such situations, a dummy event variable $\mathbf{x}$ can be used with each event having the same value of the event variable. Alternatively, instead of estimating the signal-fraction functions $p_c(\mathbf{x})$ in line 6 of algorithm 1, one can estimate the expected value of p conditional only on c . Either of these simplifications leads to an elegant solution to the problem of finding the optimal cutoffs or “working points” on ML classifier outputs.
- It is possible that the training process converges poorly in terms of fraction of events reassigned in each step, despite the estimated value of $\tilde{D}_{\text{stat}}$ saturating. It is also possible for different repetitions of the training process (with different training datasets

²⁴Near the end of the training process, it is possible for the estimated $\tilde{D}_{\text{stat}}$ value to decrease over training steps. This is caused purely by artifacts in the regression technique used in the estimation of $p_c(\mathbf{x})$. Such a decrease is otherwise not possible, as proved in appendix C.

or different initializations of η) to converge on categorizers that are substantially different from one another, despite having similar values of estimated $\tilde{D}_{\text{stat}}$.

This leeway in categorizer specification can be interpreted as certain decisions regarding category assignments of events *not mattering much*, and could be a sign that the number of categories $\mathcal{C}$ is too high (to be warranted by the amount of training data used). In such situations, it may be possible to reduce the variability in the trained categorizer by reducing the number of categories $\mathcal{C}$ without a significant drop in performance. Alternatively, one can reduce the variability in the trained classifier by sufficiently increasing the number of training events, in order to reliably probe the small variation in performance.

- Note that one has no control over the fraction of events that get placed in different categories by the trained categorizer. This could lead to very lopsided categorizations, leading to situations where increasing the number of categories $\mathcal{C}$ makes the expected number of events in some categories inconveniently small while other categories still have sufficient number of events to warrant further splitting. In such situations, one can freeze certain categories and force others to be split further by passing them through another categorizer. This is similar to the sequential splitting of events into categories used in [26].

One could also use different values of $\mathcal{C}$ in different regions of $\mathbf{x}$, with each region using a different categorizer.

- It is possible to train a categorizer stochastically instead of iteratively, by maintaining a moving estimate of $p_c(\mathbf{x})$, updated based on a few randomly sampled training events (mini-batch) in each step. Similar to stochastic k-means clustering algorithms [50, 51], such a stochastic implementation will scale up better with large training datasets and large number of categories $\mathcal{C}$.

As one can see, our training prescription is not hard-and-fast. It is more of a paradigm with several choices and variations, not all of which have been covered here, and a comprehensive exploration of these variations is beyond the scope of this work.

4 The ThickBrick package: a public implementation in Python

We provide a publicly available implementation of ThickBrick as a Python 3 [package](#) of the same name [38]. In this section we will demonstrate our ThickBrick package with a toy example. A guided tutorial of the package in the context of a physics example (a Higgs search in associate production $ZH \rightarrow (\mu^+\mu^-)(b\bar{b})$) was recently presented at the PyHEP 2020 virtual workshop [52]. Additional physics-motivated examples will be demonstrated in a forthcoming publication [53].

Toy example. For our toy example, we will use two dimensional data $\mathbf{e} \equiv (x, x_c)$ generated as follows. The background is taken to be uniformly distributed over the square-shaped

Figure 6. A 2d histogram of the signal distribution in eq. (4.2) normalized to 1 (setting $S/B = 1$).

region defined by $-10 \leq x, x_c \leq 10$.

$$b(x, x_c) = \begin{cases} B/400, & \text{if } |x|, |x_c| \leq 10 \\ 0, & \text{otherwise} \end{cases} \quad (4.1)$$

The signal is taken to be constrained within the same region with a truncated bivariate normal distribution given by

$$s(x, x_c) = \begin{cases} S A e^{-x^2/(2\sigma^2)} e^{-x_c^2/(2\sigma_c^2)}, & \text{if } |x|, |x_c| \leq 10 \\ 0, & \text{otherwise} \end{cases}, \quad (4.2)$$

with $\sigma = 2$ and $\sigma_c = 3$, and A is the normalization factor given by

$$A = \frac{2}{\pi \sigma \sigma_c} \left[\operatorname{erf} \left(\frac{10}{\sigma \sqrt{2}} \right) - \operatorname{erf} \left(\frac{-10}{\sigma \sqrt{2}} \right) \right]^{-1} \left[\operatorname{erf} \left(\frac{10}{\sigma_c \sqrt{2}} \right) - \operatorname{erf} \left(\frac{-10}{\sigma_c \sqrt{2}} \right) \right]^{-1}. \quad (4.3)$$

Figure 6 shows the unit normalized signal distribution. In this example, we will use x as the event variable $\boldsymbol{x}$ which will be used to search for the signal. x_c represents the information in $\boldsymbol{e}$ complementary to $\boldsymbol{x}$, which in a real physics example could be high dimensional. In the preprocessing step, we can take training events from the background and signal distribution in a 1:1 ratio and estimate the probability p that an event came from the signal distribution condition on (x, x_c) using a ML classifier training algorithm as per section 3.1. But since the focus of our work here is not this preprocessing step, for simplicity we will use our knowledge of the underlying distributions to directly calculate $p(x, x_c)$ as²⁵

$$p(x, x_c) = \frac{s(x, x_c)}{s(x, x_c) + b(x, x_c)}, \quad (4.4)$$

with S/B set to 1. This $p(x, x_c)$ is plotted in the top-right panel of figure 7. In this section, we will consider other values of the overall signal-to-background ratio S/B as well, but we

²⁵Note that the quality of the event selection/categorization is dependent on the quality of estimation of $p(\boldsymbol{e})$, and here we are bypassing this step.

Figure 7. Plots illustrating the toy data under consideration. The top-left panel shows the unit-normalized 2d distribution of $(x, \text{logit}(p))$. The top-right panel shows a plot of $p(x, x_c)$ given in eq. (4.4). The bottom two panels show the distribution of signal events (orange) and background events (blue) in the event variable x before any event selection. The background distribution is normalized to 1 in both, and the signal-to-background ratio is taken to be 1 (bottom-left panel, the signal is 50% of the total) and $1/9$ (bottom-right panel, the signal is 10% of the total).

will let the ThickBrick package handle the scaling internally. Therefore, in all plots in this section p will refer to the value corresponding to $S/B = 1$. The top-left panel of figure 7 shows the unit normalized 2d distribution of $(x, \text{logit}((p(x, x_c))))$. For better visualization of the spread of p at low values of $|x|$, here we use the logit function given by

$$\text{logit}(p) \equiv \ln \left(\frac{p}{1-p} \right), \quad (4.5)$$

which also has the convenient property that scaling the signal cross-section or, equivalently, changing S/B , only changes $\text{logit}(p)$ by an additive constant. In the top-left panel of figure 7, the function $\text{logit}(p)$ appears to be constrained from below only because of the truncation of the data at the boundaries $x_c = \pm 10$. Finally, the bottom two panels of figure 7 show the distribution of signal events (orange) and background events (blue) in the event variable x before any event selection. The bottom-left panel has $S/B = 1$ (the signal is 50% of the total) and the bottom-right panel has $S/B = 1/9$ (the signal is 10% of the total), with the background distributions in both panels normalized to 1.

Event selection. Let us start off with event selection (`C = 1`, `allow_rejection = True`) with $\tilde{D}_{\text{Pear}\chi^2}$ as the distance measure (so that we do not have to consider the effect of

Figure 8. *Best flat-threshold selector.* The black horizontal line in the top-left panel shows the constant selection threshold that maximizes $\tilde{D}_{\text{Pear}\chi^2}$. The region bounded by the orange curve in the top-right panel is the region in (x, x_c) selected by this threshold. The colors in the plot have the same meaning as in the top-right plot of figure 7. The bottom two panels show the signal and background distributions in x after event selection. The background distribution *before selection* is normalized to 1 in both, and the signal-to-background ratio *before selection* is taken to be 1 (bottom-left panel) and 1/9 (bottom-right panel). This allows for direct comparisons with the plots in figure 7.

changing S/B). Figure 8 illustrates the best event selector with an x -independent threshold on p . The black horizontal line in the top-left panel shows the constant threshold that maximizes $\tilde{D}_{\text{Pear}\chi^2}$ — events above this threshold will be selected. The top-right panel shows the region in the full phase-space that is selected by this threshold (the region bounded by the orange curve). The bottom two panels show the signal and background distributions in x after event selection. The background distribution *before selection* is normalized to 1 in both, and the signal-to-background ratio *before selection* is taken to be 1 (bottom-left panel) and 1/9 (bottom-right panel). These plots can therefore be directly compared to the bottom panels of figure 7. Notice the strong correlation between the selection criterion and the event variable x in the top-right plots, and the resulting background shaping in the bottom panels.

Next, let us consider a selector trained as per our prescription to maximize the same distance measure $\tilde{D}_{\text{Pear}\chi^2}$. We used the 1d Nadaraya-Watson kernel regression technique [54, 55] to estimate p_c in each iterative step (implemented within the ThickBrick package). A training sample of 50000 background and signal events each was used to

Figure 9. Event selector trained by ThickBrick to maximize $\tilde{D}_{\text{Pear}\chi^2}$ with S/B set to 1. All panels are analogous to their counterparts in figure 8. The region bounded by the orange horizontal lines in the top-right panel is selected.

train the selector, and the training converged in 10 training steps. Figure 9 illustrates the trained event selector, and all the panels are analogous to their counterparts in figure 8 — in the top-right panel, the region within the two horizontal orange lines is selected. Note how the selection is decorrelated from the event variable, confirming that event selection is occurring only based on the complementary information contained in x_c .

A word on decorrelation. In this toy example x_c and x were constructed to be independent of each other under both the signal and background distributions for illustrative purposes only — in our training process, we never use this fact. The optimal selection thresholds were learned only from the distribution of events in the (x, p) space. One can see that our event selector uses only complementary information from the fact that the selection cuts in the top-right panel of figure 9 are horizontal.

Note that different distance measures will decorrelate the selection criteria from the event variable $\mathbf{x}$ to different degrees, and $\tilde{D}_{\text{Pear}\chi^2}$ is a special distance measure that leads to categorizers that are *completely* decorrelated from the event variable $\mathbf{x}$. Since the decorrelation properties of our distance measures are not the main focus of this paper, we only briefly discuss this point in section 6. In figure 10, we illustrate a selector trained to maximize $\tilde{D}_{\text{Neym}\chi^2}$, which is at the other end of the spectrum in terms of degree of decorrelation. Since training with distances other than $\tilde{D}_{\text{Pear}\chi^2}$ is sensitive to the value of S/B , and the contrast between different distance measures increases with increasing S/B , we will use a high signal-to-background ratio of 1 for illustrative purposes.

Figure 10. Event selector trained by ThickBrick to maximize $\tilde{D}_{\text{Neym}\chi^2}$. The left and right panels are analogous to the top-right and bottom-left panels of figure 9, respectively.

Figure 11. The same as the top two panels in figure 9, but for the case of event categorizers trained by ThickBrick. Event rejection was turned off and the number of categories was two (three) in the top (bottom) row. In the top-right panel, the region within the orange horizontal lines comprises category 2, while the two regions outside the orange lines comprise category 1. In the bottom-right panel, the region within the orange horizontal lines comprises category 3, the two regions bounded by an orange and a white line comprise category 2, and the two regions outside the white lines comprise category 1.

Note that the implication of figure 10 is *not* that the corresponding event selector bases its decision partially on information non-complementary to $\mathbf{x}$. In our toy example, the threshold on p at a given value of x will not change if we scale up the number of

signal events at other values of x . The correct interpretation is that the selector bases its decision on the *number* density distributions of x_c under $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$ conditional on x , instead of the *probability* distributions of x_c conditional on x — the background and signal distributions in x affect the selection criteria through this subtle difference.

Event categorization. Next we will consider event categorization with event rejection turned off (`allow_rejection = False`). As before, we will only consider $\tilde{D}_{\text{Pear}\chi^2}$ and use a training sample of 50000 background and signal events each. The top and bottom rows of figure 11 illustrate categorizers trained with number of categories set to 2 and 3, respectively. In both cases, the black curves in the plots on left show the threshold between categories. The right panels show the corresponding boundaries in the (x, x_c) phase-space. In the top-right panel, the region within the orange horizontal lines comprises category 2, while the two regions outside the orange lines comprise category 1. In the bottom-right panel, the region within the orange horizontal lines comprises category 3, the two regions bounded by an orange and a white line comprise category 2, and the two regions outside the white lines comprise category 1.

Performance comparison using $\tilde{D}_{\text{Pear}\chi^2}$. Figure 12 illustrates the performance of the event selectors and categorizers trained using our prescription, as captured by the estimated $\tilde{D}_{\text{Pear}\chi^2}$ with S/B set to 1/9. For comparison, the best event selector that uses a flat threshold on $p(e)$ is also shown, along with the theoretical upper-limit on $D_{\text{Pear}\chi^2}$ achievable by analyzing the distribution of events in the full phase-space (x, x_c) . The x -axis of the plot shows the number of selected categories $\mathcal{C}$. The red cross corresponds to the best flat-threshold selector, while the blue triangle corresponds to an event selector trained using our prescription. The three black dots all correspond to event categorizers with different values of $\mathcal{C}$ (all with `allow_rejection = True`). The horizontal green dashed line in figure 12 shows the theoretical upper-limit on $\tilde{D}_{\text{Pear}\chi^2}$ for $\mathcal{C} \rightarrow \infty$, which is set by the data processing inequality. There are no appreciable error bars to report in figure 12. Note that, as expected, the event selector trained with ThickBrick performs better than a flat-threshold selector in terms of the statistical distance measure, $\tilde{D}_{\text{Pear}\chi^2}$. Also note the diminishing returns with increasing values of $\mathcal{C}$, as anticipated in section 3.2.

Performance comparison using the median significance for signal discovery. In order to get a median significance value for discovering a signal, we re-write D_{KL} as

$$D_{\text{KL}} = \ln \left[\frac{N_{\text{sel}}^{N_{\text{sel}}} e^{-N_{\text{sel}}}}{B_{\text{sel}}^{N_{\text{sel}}} e^{-B_{\text{sel}}}} \right] + \sum_{c=1}^{\mathcal{C}} \int_{\Omega_x} d\mathbf{x} \left[n_c(\mathbf{x}) \ln \left[\frac{n_c(\mathbf{x})/N_{\text{sel}}}{b_c(\mathbf{x})/B_{\text{sel}}} \right] \right], \quad (4.6)$$

where B_{sel} and N_{sel} are the expected number of non-rejected events under $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$, respectively:

$$B_{\text{sel}} = \sum_{c=1}^{\mathcal{C}} B_c, \quad N_{\text{sel}} = \sum_{c=1}^{\mathcal{C}} N_c. \quad (4.7)$$

This shows that D_{KL} is the value of the test statistic $\ln(\mathcal{P}_{b+s}/\mathcal{P}_b)$ for an “Asimov dataset” [42] produced under $\mathbb{H}_{b+s}$, where $\mathcal{P}_{b+s}$ and $\mathcal{P}_b$ are the probability densities for

Figure 12. Comparison of different event selectors and categorizers in terms of the estimated $\tilde{D}_{\text{Pear}\chi^2}$ with $S/B = 1/9$. The red cross corresponds to the best flat-threshold selector, while the blue triangle corresponds to an event selector trained using ThickBrick. The three black dots all correspond to event categorizers with different values of $\mathcal{C}$ (all with `allow_rejection = True`). The horizontal green dashed line shows the theoretical upper-limit on $\tilde{D}_{\text{Pear}\chi^2}$ for $\mathcal{C} \rightarrow \infty$.

observing the dataset (with events characterized by the category they belong to and their $\mathbf{x}$ values) under $\mathbb{H}_{b+s}$ and $\mathbb{H}_b$, respectively. Using the results of [42], for large N , the median significance of discovering a signal, given $\mathbb{H}_{b+s}$ is true, is given by

$$Z_{\text{median}} = \sqrt{2 \ln \left(\frac{\mathcal{P}_{b+s}}{\mathcal{P}_b} \right)} \text{ for Asimov dataset} = \sqrt{2 D_{\text{KL}}} = \sqrt{N \tilde{D}_{\text{KL}}} . \quad (4.8)$$

For the event selectors and categorizers shown in figure 12, the median significance of discovery (estimated using (4.8)) is shown in figure 13, for $S/B = 1/99$ and $N = 50000$. The green line in figure 13 corresponds to the median significance achievable by performing the likelihood-ratio test on the full event information $\mathbf{e}$.

5 Conclusions

In this paper we have provided an optimal event selection and categorization prescription called ThickBrick to maximize the statistical significance of a potential signal excess (in the distribution of a low dimensional event variable). This statistical significance is captured by one of six different statistical distance measures given in eq. (2.1).

Our prescription calls for first learning the probability $p(\mathbf{e})$ that an event (sampled from the “sig+bg” hypothesis, $\mathbb{H}_{b+s}$) is a signal event, possibly using machine learning methods. Let us refer to the probability thus learned as “the machine learning output”. ThickBrick then finds the optimal selection thresholds (for event selectors) or boundary between categories (for event categorizers) in the machine learning output in an iterative manner.²⁶ These thresholds and boundaries are dependent on the event variable. The

²⁶An alternative approach to finding these optimal thresholds may be possible using BDTs (trained to maximize our performance measures) with $\mathbf{x}(\mathbf{e})$ and $p(\mathbf{e})$ as inputs.

Figure 13. Comparison of different event selectors and categorizers in terms of the median significance Z_{median} , for $S/B = 1/99$ and $N = 50000$. The event selectors and categorizers depicted here are the same as the ones in figure 12. The horizontal green dashed line shows the median significance achievable by performing the likelihood-ratio test on the full event information $\mathbf{e}$.

thresholds can be thought of as chosen to extract the best sensitivity out of each $\mathbf{x}$ -“bin”.

Note that since $p(\mathbf{e})$ contains all the relevant useful information, in principle it can be used directly for statistical inference, for example, by constructing the Neyman-Pearson statistic. However, in practice this is complicated because $p(\mathbf{e})$ is estimated from simulations which may have known and unknown systematic uncertainties. In such cases, analyses based on event variables tend to be more robust because they allow for extensive simulation validation, e.g., using control region validation. Such analyses still benefit from the information orthogonal to the event variable $\mathbf{x}$ by using that information for event selection and classification. The technique introduced in this paper optimizes this aspect of the analysis and demands only very minimal changes to the current practices — event selectors and categorizers trained by our technique can be used as drop-in replacements for the currently used ones. Any errors and uncertainties in the estimation of $p(\mathbf{e})$ will not affect the robustness of the analyses, only the optimality of event selection.

Event selectors (categorizers) trained using ThickBrick are expected to outperform typical ML-classifier-based selectors (categorizers) in terms of statistical sensitivity, as demonstrated using a toy example in section 4. In addition, our event selectors and categorizers will also be decorrelated from the event variable used in the analysis, which can help reduce the impact of systematic uncertainties on signal search analyses.

6 Epilogue: decorrelated classifiers and taggers

The distance measure $D_{\text{Pear}\chi^2}$ from eq. (2.1b) is given by

$$D_{\text{Pear}\chi^2} = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \frac{s_c^2(\mathbf{x})}{b_c(\mathbf{x})} . \quad (6.1)$$

As noted in section 3.2, scaling the signal number density by *an overall factor* ($s(\mathbf{e}) \rightarrow \lambda s(\mathbf{e})$) only changes $D_{\text{Pear}\chi^2}$ by a multiplicative factor ($D_{\text{Pear}\chi^2} \rightarrow \lambda^2 D_{\text{Pear}\chi^2}$). As a result, the ordering of categorizers by their $\tilde{D}_{\text{stat}}$ values is independent of S/B , and hence, so is the optimal categorizer.

What is interesting is that the optimality of the categorizer that maximizes $\tilde{D}_{\text{Pear}\chi^2}$ is also unaffected by an $\mathbf{x}$ dependent scaling given by

$$s(\mathbf{e}) \rightarrow \lambda(\mathbf{x}(\mathbf{e})) s(\mathbf{e}) . \quad (6.2)$$

This is because at a given value of $\mathbf{x}$, the transformation in eq. (6.2) only changes $\sum_c s_c^2(\mathbf{x})/b_c(\mathbf{x})$ by the multiplicative factor $\lambda^2(\mathbf{x})$, and the optimal categorizer maximizes $\sum_c s_c^2/b_c$ at each value of $\mathbf{x}$.

This means that the categorizer which maximizes $\tilde{D}_{\text{Pear}\chi^2}$ will be *completely* decorrelated from $\mathbf{x}$ — it will not “learn” from the distributions of signal and background over $\mathbf{x}$ in the training data, and will only learn from their respective probability distributions conditional on $\mathbf{x}$ (see figures 9 and 11). This property is often desirable, especially when the distribution over $\mathbf{x}$ (of signal and/or background) is uncertain [31, 35, 56]. For example, this decorrelation can be useful in jet substructure based boosted-decay taggers (decorrelation with jet mass is desirable). Another potential application is in decorrelating quark and gluon jet classifiers trained on “naturally mixed samples” (like quark enriched Z+jet samples and gluon enriched dijet samples) [57–59] from observables like jet energy and pseudorapidity, which are controlled by the different underlying kinematics of the corresponding mixed samples.

While $D_{\text{Pear}\chi^2}$ is a good performance measure from the standpoint of signal discovery, there may be other performance measures better suited to other applications, which share the same decorrelation property. For example, for a classification problem which places signal and background on the same footing, it may be better to use D_{classify} , defined below, as the performance measure to maximize (iteratively).

$$D_{\text{classify}} = - \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \sqrt{s_c(\mathbf{x})b_c(\mathbf{x})}. \quad (6.3)$$

This form of D_{classify} has connections to the Bhattacharyya coefficient [44].

In fact, we can generalize $D_{\text{Pear}\chi^2}$ and D_{classify} to a family of performance measures parameterized by ν , all of which share the decorrelation property,²⁷ as follows:

$$D_{\text{classify},\nu} = \text{sgn}[\nu(\nu-1)] \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} s_c^\nu(\mathbf{x}) b_c^{1-\nu}(\mathbf{x}) , \quad \text{for } \nu(\nu-1) \neq 0 . \quad (6.4)$$

Similarly, we can also generalize D_{classify} to multi-class classification problems. We will explore these ideas in a future publication.

²⁷Although the utility of such a generalization is presently unclear.

Acknowledgments

The work of KM and PS is supported in part by the United States Department of Energy under Grant No. DE-SC0010296. PS is grateful to the LHC Physics Center at Fermilab for hospitality and financial support as part of the Guests and Visitors Program in the summer of 2019. PS would like to thank Prof. Arunava Banerjee whose course on machine learning at the University of Florida jump-started this project. The authors would like to thank the following colleagues for useful discussions and inputs: Darin Acosta, Felix Kling, Jacobo Konigsberg, Danny Noonan, Andrzej Novak, Brendan O’Brien, Nathaniel Joseph Pastika, Jake Rosenzweig, Reinhard Schwienhorst, and Sean-Jiun Wang. The authors would also like to thank the anonymous referees of this paper for their detailed comments.

Data availability. The code and data that support the findings of this study are openly available at the following URL: <https://gitlab.com/prasanthcakewalk/code-and-data-availability/> under the directory named `arXiv_1911.12299`.

A Statistical distances

In this section, we will derive the expressions for the different statistical distances given in eq. (2.1). Let $\mathcal{S}$ be an ordered²⁸ collection of collider events. A given instance of $\mathcal{S}$ will be treated as one of the possible outcomes of the experiment performed to produce it. The events in $\mathcal{S}$ are independent and identically distributed. The total number of events in $\mathcal{S}$ is a Poisson distributed random variable whose mean could possibly depend on the underlying theory.

After categorizing events in $\mathcal{S}$ into C categories (and possibly rejecting some), let $\mathcal{S}_c$ be the ordered²⁸ collection of events in category c . Let $\Upsilon_c \equiv [k_c, (\mathbf{x}_{c,1}, \mathbf{x}_{c,2}, \dots, \mathbf{x}_{c,k_c})]$ represent the list of $\mathbf{x}$ values for events in category c , where k_c is the number of events in category c .

Since hypothesis testing is to be performed by analyzing the Υ -s, we will next write down their probability densities under the two hypotheses. Let $\mathcal{P}_b(\cdot) \equiv \mathcal{P}(\cdot; \mathbb{H}_b)$ and $\mathcal{P}_{b+s}(\cdot) \equiv \mathcal{P}(\cdot; \mathbb{H}_{b+s})$ represent the probability density of $\cdot$ under the null and alternative hypotheses respectively. We can write the probability densities of Υ_c as

$$\mathcal{P}_b(\Upsilon_c) = \frac{e^{-B_c} B_c^{k_c}}{k_c!} \prod_{i=1}^{k_c} \frac{b_c(\mathbf{x}_{c,i})}{B_c} , \quad (\text{A.1a})$$

$$\mathcal{P}_{b+s}(\Upsilon_c) = \frac{e^{-N_c} N_c^{k_c}}{k_c!} \prod_{i=1}^{k_c} \frac{n_c(\mathbf{x}_{c,i})}{N_c} . \quad (\text{A.1b})$$

These probability densities are normalized as²⁹

$$\int d\Upsilon_c \mathcal{P}(\Upsilon_c) \equiv \sum_{k_c=0}^{\infty} \int_{\Omega_{\mathbf{x}}} d\mathbf{x}_{c,1} \dots \int_{\Omega_{\mathbf{x}}} d\mathbf{x}_{c,k_c} \mathcal{P}(\Upsilon_c) = 1 . \quad (\text{A.2})$$

²⁸Ordering of events can be derived from some form of event id. Ordering is demanded just so we don’t have to worry about combinatorial factors when writing down the probability (density) of a particular instance of $\mathcal{S}$, $\mathcal{S}_c$, or Υ_c , since the events are distinguishable.

²⁹Eq. (A.2) implicitly specifies the reference measure with respect to which the probability densities of Υ_c and $\{\Upsilon\}$ are defined.

Since the different Υ -s are independent of each other, their joint probability density is given by

$$\mathcal{P}_b(\{\Upsilon\}) \equiv \mathcal{P}_b(\Upsilon_1, \dots, \Upsilon_C) = \prod_{c=1}^C \mathcal{P}_b(\Upsilon_c) \quad , \quad (\text{A.3a})$$

$$\mathcal{P}_{b+s}(\{\Upsilon\}) \equiv \mathcal{P}_{b+s}(\Upsilon_1, \dots, \Upsilon_C) = \prod_{c=1}^C \mathcal{P}_{b+s}(\Upsilon_c) \quad . \quad (\text{A.3b})$$

From eqs. (A.1a), (A.1b), (A.3a), and (A.3b), the ratio of the probability densities can be written as

$$\frac{\mathcal{P}_{b+s}(\{\Upsilon\})}{\mathcal{P}_b(\{\Upsilon\})} = \prod_{c=1}^C \left[e^{-N_c+B_c} \prod_{i=1}^{k_c} \frac{n_c(\mathbf{x}_{c,i})}{b_c(\mathbf{x}_{c,i})} \right] \quad . \quad (\text{A.4})$$

This expression will be used in the derivation of some of our statistical distances.

A.1 $D_{\text{Neym}\chi^2}$ and $D_{\text{Pear}\chi^2}$: asymptotic expected values of Neyman- χ^2 and Pearson- χ^2 statistics

For a binned analysis of the distribution of $\mathbf{x}$ within each category, the Neyman- χ^2 and Pearson- χ^2 test statistics are given by

$$\text{Neym}\chi^2 = \sum_{c=1}^C \sum_{i \in \mathbf{x} \text{ bins}} \frac{(\text{obs}_{c,i} - \text{exp}_{c,i})^2}{\text{obs}_{c,i}} \quad , \quad (\text{A.5a})$$

$$\text{Pear}\chi^2 = \sum_{c=1}^C \sum_{i \in \mathbf{x} \text{ bins}} \frac{(\text{obs}_{c,i} - \text{exp}_{c,i})^2}{\text{exp}_{c,i}} \quad . \quad (\text{A.5b})$$

where $\text{exp}_{c,i}$ and $\text{obs}_{c,i}$ are respectively the expected and observed number of events in bin i of category c . The summation is only performed over bins with non-zero $\text{obs}_{c,i}$ for Neyman- χ^2 , and non-zero $\text{exp}_{c,i}$ for Pearson- χ^2 .

Let $N_{c,i}$, $B_{c,i}$, and $S_{c,i}$ be the number densities $n_c(\mathbf{x})$, $b_c(\mathbf{x})$, and $s_c(\mathbf{x})$ integrated over the bin i in category c . For both variants of the χ^2 statistic, let χ_b^2 and χ_{b+s}^2 be the value of the statistic calculated with $\text{exp}_{c,i}$ set to $B_{c,i}$ and $N_{c,i}$ respectively.

In the asymptotic and fine binning limit we can show, for both variants of the χ^2 statistic, that

$$\mathbb{E} \left[\chi_b^2 ; \mathbb{H}_{b+s} \right] \rightarrow D_{\chi^2} \quad , \quad (\text{A.6a})$$

$$\left[\mathbb{E} \left[\chi_b^2 ; \mathbb{H}_{b+s} \right] - \mathbb{E} \left[\chi_b^2 ; \mathbb{H}_b \right] \right] \rightarrow D_{\chi^2} \quad , \quad (\text{A.6b})$$

$$\mathbb{E} \left[\chi_b^2 - \chi_{b+s}^2 ; \mathbb{H}_{b+s} \right] \rightarrow D_{\chi^2} \quad , \quad (\text{A.6c})$$

where $\mathbb{E}[\cdot ; \mathbb{H}_b]$ and $\mathbb{E}[\cdot ; \mathbb{H}_{b+s}]$ represent the expected value of $\cdot$ under $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$ respectively, and the two variants of D_{χ^2} are the $D_{\text{Neym}\chi^2}$ and $D_{\text{Pear}\chi^2}$ defined in equations (2.1a) and (2.1b).

The left-hand side of eq. (A.6a) is a measure of how poorly $\mathbb{H}_b$ would fit data produced under $\mathbb{H}_{b+s}$. The left-hand side of eq. (A.6b) is a measure of how much worse $\mathbb{H}_b$ would fit data produced under $\mathbb{H}_{b+s}$ than it would fit data produced under $\mathbb{H}_b$. And the left-hand

side of eq. (A.6c) is a measure of how much worse $\mathbb{H}_b$ would be than $\mathbb{H}_{b+s}$ at fitting data produced under $\mathbb{H}_{b+s}$.

We will omit an explicit proof of statements (A.6a)–(A.6c) here, and instead only provide an outline. We begin with the following observations: if X is a Poisson distributed random variable with mean μ , then the expected value of X^2 is $\mu^2 + \mu$. Similarly, the expected value of $1/X$ for positive X tends to $1/\mu + 1/\mu^2$ in the large μ limit [62], and the probability of X being 0 vanishes in the large μ limit. Using these ideas, we can show that asymptotically (up to the fastest growing term), the left-hand sides of eqns (A.6a)–(A.6c) tend to

$$\sum_{c=1}^C \sum_{i \in \mathbf{x} \text{ bins}} \frac{(N_{c,i} - B_{c,i})^2}{N_{c,i}} \quad (\text{A.7})$$

for Neyman- χ^2 and

$$\sum_{c=1}^C \sum_{i \in \mathbf{x} \text{ bins}} \frac{(N_{c,i} - B_{c,i})^2}{B_{c,i}} \quad (\text{A.8})$$

for Pearson- χ^2 . These are approximations to the expressions for $D_{\text{Neym}\chi^2}$ and $D_{\text{Neym}\chi^2}$ in eqs.eq. (2.1a) and eq. (2.1b), and the approximations can be made arbitrarily precise in the fine binning limit.

A.2 D_{KL} : Kullback-Leibler divergence from null to alternative

The Kullback-Leibler divergence from the null hypothesis to the alternative hypothesis (for the Υ -s) is given by

$$D_{\text{KL}} \equiv D_{\text{KL}}(\mathbb{H}_{b+s} \parallel \mathbb{H}_b) = \int d\{\Upsilon\} \mathcal{P}_{b+s}(\{\Upsilon\}) \ln \left[\frac{\mathcal{P}_{b+s}(\{\Upsilon\})}{\mathcal{P}_b(\{\Upsilon\})} \right] , \quad (\text{A.9})$$

where $\int d\{\Upsilon\}$ is a shorthand for

$$\int d\{\Upsilon\} \equiv \prod_{c=1}^C \left[\sum_{k_c=0}^{\infty} \int_{\Omega_{\mathbf{x}}} d\mathbf{x}_{c,1} \dots \int_{\Omega_{\mathbf{x}}} d\mathbf{x}_{c,k_c} \right] . \quad (\text{A.10})$$

Plugging eq. (A.4) into eq. (A.9), and using $s_c = n_c - b_c$ and $S_c = N_c - B_c$, we get

$$D_{\text{KL}} = \int d\{\Upsilon\} \mathcal{P}_{b+s}(\{\Upsilon\}) \sum_{c=1}^C \left[-S_c - \sum_{i=1}^{k_c} \ln \left[1 - \frac{s_c(\mathbf{x}_{c,i})}{n_c(\mathbf{x}_{c,i})} \right] \right] . \quad (\text{A.11})$$

Using the fact that the different Υ -s are independent of each other, and the fact that $\mathbf{x}_{c,i}$ for different i values are independent of each other, we can simplify this to

$$D_{\text{KL}} = \sum_{c=1}^C \left[-S_c - \left[\sum_{k_c=0}^{\infty} \frac{e^{-N_c} N_c^{k_c}}{k_c!} k_c \right] \frac{1}{N_c} \int_{\Omega_{\mathbf{x}}} d\mathbf{x} n_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] \quad (\text{A.12})$$

$$= \sum_{c=1}^C \left[-S_c - \int_{\Omega_{\mathbf{x}}} d\mathbf{x} n_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] \quad (\text{A.13})$$

$$= \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \left[-s_c(\mathbf{x}) - n_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] . \quad (\text{A.14})$$

The last equation matches the expression for D_{KL} given in eq. (2.1c).

A.3 D_{revKL} : Kullback-Leibler divergence from alternative to null

The Kullback-Leibler divergence from the alternative hypothesis to the null hypothesis (for the Υ -s) is given by

$$D_{\text{revKL}} \equiv D_{\text{KL}}(\mathbb{H}_b \parallel \mathbb{H}_{b+s}) = \int d\{\Upsilon\} \mathcal{P}_b(\{\Upsilon\}) \ln \left[\frac{\mathcal{P}_b(\{\Upsilon\})}{\mathcal{P}_{b+s}(\{\Upsilon\})} \right] . \quad (\text{A.15})$$

Using manipulations similar to the ones we did for D_{KL} , we can show that

$$D_{\text{revKL}} = \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \left[s_c(\mathbf{x}) + b_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] , \quad (\text{A.16})$$

which matches the expression for D_{revKL} given in eq. (2.1d). This expression can also be derived from eq. (A.14) by making the substitutions $n_c \longleftrightarrow b_c$ and $s_c \rightarrow -s_c$.

A.4 D_{J} : Jeffreys divergence

Jeffreys divergence is a symmetrized version of Kullback-Leibler divergence given by the sum of the two asymmetric KL divergences. From eqs. (A.14) and (A.16) we can see that

$$D_{\text{J}} = D_{\text{KL}} + D_{\text{revKL}} \quad (\text{A.17})$$

$$= \sum_{c=1}^C \int_{\Omega_{\mathbf{x}}} d\mathbf{x} \left[-s_c(\mathbf{x}) \ln \left[1 - \frac{s_c(\mathbf{x})}{n_c(\mathbf{x})} \right] \right] . \quad (\text{A.18})$$

This matches the expression for D_{J} given in eq. (2.1e).

A.5 D_{B} : Bhattacharyya distance

The Bhattacharyya distance between null and alternative hypotheses (for the Υ -s) is given by

$$D_{\text{B}} = -\ln \left[\int d\{\Upsilon\} \sqrt{\mathcal{P}_b(\{\Upsilon\}) \mathcal{P}_{b+s}(\{\Upsilon\})} \right] . \quad (\text{A.19})$$

We can rewrite this as

$$e^{-D_{\text{B}}} = \int d\{\Upsilon\} \mathcal{P}_{b+s}(\{\Upsilon\}) \sqrt{\frac{\mathcal{P}_b(\{\Upsilon\})}{\mathcal{P}_{b+s}(\{\Upsilon\})}} . \quad (\text{A.20})$$

Plugging eq. (A.4) into eq. (A.20) we get

$$e^{-D_{\text{B}}} = \left[\prod_{c=1}^C e^{S_c/2} \right] \left[\int d\{\Upsilon\} \mathcal{P}_{b+s}(\{\Upsilon\}) \prod_{c=1}^C \prod_{i=1}^{k_c} \sqrt{\frac{b_c(\mathbf{x}_{c,i})}{n_c(\mathbf{x}_{c,i})}} \right] . \quad (\text{A.21})$$

Plugging in eqs. (A.1b), (A.3b), and (A.10) here and using the fact that the different Υ -s are independent of each other, and the fact that $\mathbf{x}_{c,i}$ for different i values are independent of each other, we can simplify this to

$$e^{-D_{\text{B}}} = \prod_{c=1}^C \left[e^{S_c/2} e^{-N_c} \sum_{k_c=0}^{\infty} \frac{1}{k_c!} \left[\int_{\Omega_{\mathbf{x}}} d\mathbf{x} n_c(\mathbf{x}) \sqrt{\frac{b_c(\mathbf{x})}{n_c(\mathbf{x})}} \right]^{k_c} \right] . \quad (\text{A.22})$$

Using the power series representation of the exponential function given by

$$e^A = \sum_{k=0}^{\infty} \frac{A^k}{k!} , \quad (\text{A.23})$$

we can rewrite eq. (A.22) as

$$e^{-D_B} = \prod_{c=1}^C \left[e^{S_c/2} e^{-N_c} \exp \left[\int_{\Omega_x} d\mathbf{x} \, n_c(\mathbf{x}) \sqrt{\frac{b_c(\mathbf{x})}{n_c(\mathbf{x})}} \right] \right] \quad (\text{A.24})$$

$$= \exp \left[\sum_{c=1}^C \int_{\Omega_x} d\mathbf{x} \left[\frac{s_c(\mathbf{x})}{2} - n_c(\mathbf{x}) + n_c(\mathbf{x}) \sqrt{\frac{b_c(\mathbf{x})}{n_c(\mathbf{x})}} \right] \right] . \quad (\text{A.25})$$

This matches the expression for D_B given in eq. (2.1f) (after taking natural-log and multiplying by -1).

B An interpretation of Kullback-Leibler divergence

Let us consider a generic statistical test that decides between $\mathbb{H}_b$ or $\mathbb{H}_{b+s}$, based on data (post-categorization, i.e., after reduction of events to their c and $\mathbf{x}$ values).

Power of the test, $1 - \beta$, is the probability of the test choosing $\mathbb{H}_{b+s}$, given that $\mathbb{H}_{b+s}$ is true. It captures the ability of the test to detect the presence of a signal.

Size (or significance level) of the test, α , is the probability of the test choosing $\mathbb{H}_{b+s}$, given that $\mathbb{H}_b$ is true. It captures the false alarm rate, and is typically set to $\sim 3 \times 10^{-7}$ (5σ) in high energy physics.

A statistical test that has the greatest power $1 - \beta^*$ among all tests of a given size α^* is called a uniformly most powerful test or a UMP test (with α^* being a tunable parameter in the test). The likelihood-ratio test is a UMP for testing point hypotheses (like our $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$). It can be shown for point hypotheses that

1. For a UMP test of a *given (fixed) power* $0 \leq 1 - \beta^* < 1$,

$$\alpha^* \longrightarrow e^{-D_{\text{KL}} + o(D_{\text{KL}})} \quad \text{asymptotically} \quad (\text{B.1})$$

2. For a UMP test of a *given (fixed) size* $0 < \alpha^* \leq 1$,

$$1 - \beta^* \longrightarrow 1 - e^{-D_{\text{revKL}} + o(D_{\text{revKL}})} \quad \text{asymptotically} \quad (\text{B.2})$$

where “asymptotically” refers to the large number of events limit. $o(D_{\text{KL}})$ and $o(D_{\text{revKL}})$ are functions asymptotically dominated by D_{KL} and D_{revKL} , respectively. The formal definition of the little- o notation is:

$$f = o(g) \quad \implies \quad \frac{|f|}{g} \longrightarrow 0 \quad \text{asymptotically.} \quad (\text{B.3})$$

Note that D_{KL} and D_{revKL} both scale linearly with the integrated luminosity. This means that any non-zero difference between D_{KL} and D_{revKL} will not be asymptotically dominated by either of them.

When the goal of event categorization is to maximize the chance of signal discovery (assuming its existence), we want to maximize the power for a given size (say, p-value corresponding to 5σ). This can be done by maximizing D_{revKL} , according to statement (B.2). When there is no a priori expectation for the size of the signal cross-section, it makes sense to work with a signal cross-section near the low end of discoverability.

On the other hand, when the goal of event categorization is to maximize the chance of ruling out the existence of a signal (assuming its absence), we want to minimize the size for a given power (say, 95%: this is the confidence that the signal does not exist — otherwise it would have been seen). This can be done by maximizing D_{KL} , according to statement (B.1). When there is no a priori expectation for the size of the signal cross-section and the goal is to set the best upper-limit on it, it makes sense to work with a signal cross-section value near the expected upper-limit.

This discussion is purely theoretical and it is likely that the difference between training with $\tilde{D}_{\text{revKL}}$ and $\tilde{D}_{\text{KL}}$ will not be practically significant in most situations. Nevertheless, we find this (slight) misalignment between the goals of event categorization for signal discovery and upper-limit setting a curiosity worth noting.

Proof of statement (B.1): for any hypothesis test performed using the categorized events described by $\{\Upsilon\}$, the power of the test $1 - \beta$ and the size of the test α are given by

$$1 - \beta = \int_{\text{critical region}} d\{\Upsilon\} \mathcal{P}_{\text{b+s}}(\{\Upsilon\}), \quad (\text{B.4})$$

$$\alpha = \int_{\text{critical region}} d\{\Upsilon\} \mathcal{P}_{\text{b}}(\{\Upsilon\}). \quad (\text{B.5})$$

where the integrations are performed over the critical region of the test. From these, we can write

$$\alpha = \int_{\text{critical region}} d\{\Upsilon\} \mathcal{P}_{\text{b+s}}(\{\Upsilon\}) \frac{\mathcal{P}_{\text{b}}(\{\Upsilon\})}{\mathcal{P}_{\text{b+s}}(\{\Upsilon\})} \quad (\text{B.6})$$

$$= (1 - \beta) \mathbb{E}_{\text{b+s}} \left[\frac{\mathcal{P}_{\text{b}}(\{\Upsilon\})}{\mathcal{P}_{\text{b+s}}(\{\Upsilon\})} \mid \{\Upsilon\} \in \text{critical region} \right], \quad (\text{B.7})$$

where $\mathbb{E}_{\text{b+s}}$ represents the expectation under $\mathbb{H}_{\text{b+s}}$. This result holds for any hypothesis test.

Now, let us turn to a UMP test performed using the following test statistic T .

$$T(\{\Upsilon\}) \equiv \ln \left[\frac{\mathcal{P}_{\text{b}}(\{\Upsilon\})}{\mathcal{P}_{\text{b+1}}(\{\Upsilon\})} \right]. \quad (\text{B.8})$$

The critical region of the test contains all the outcomes $\{\Upsilon\}$ for which $T(\{\Upsilon\})$ is *below a certain threshold*. We are interested in the asymptotic distribution of the test statistic T under $\mathbb{H}_{\text{b+s}}$. From (A.4), we can write T as

$$T(\{\Upsilon\}) = \sum_{c=1}^C \left[(N_c - B_c) + \sum_{i=1}^{k_c} \ln \left[\frac{b_c(\mathbf{x}_{c,i})}{n_c(\mathbf{x}_{c,i})} \right] \right] \quad (\text{B.9})$$

$$= \sum_{c=1}^C (N_c - B_c) + \sum_{c=1}^C \sum_{i=1}^{k_c} \ln \left[\frac{b_c(\mathbf{x}_{c,i})}{n_c(\mathbf{x}_{c,i})} \right] \quad (\text{B.10})$$

This shows that T can be written as the sum of $\ln(b_c/n_c)$ over all the selected events plus a constant term (independent of the observed $\{\Upsilon\}$, but dependent on the integrated luminosity). In the asymptotic limit (large expected number of events), this sum converges in distribution to the normal distribution, as per the central limit theorem.³⁰ Let μ and σ^2 be the mean and variance of T . Two keys points here are,

- Asymptotically, both μ and σ^2 scale linearly with the integrated luminosity of the experiment.
- Under the alternative hypothesis $\mathbb{H}_{b+s}$, the expected value of T is $-D_{\text{KL}}$, as can be seen from (A.9). In other words, $\mu = -D_{\text{KL}}$.

Let T^* be the threshold value of T that leads to a given power $1 - \beta^* \in [0, 1)$. The critical region of the test is $T(\{\Upsilon\}) \leq T^*$. Asymptotically, we have

$$\Phi\left(\frac{T^* - \mu}{\sigma}\right) \rightarrow 1 - \beta^* . \quad (\text{B.11})$$

where Φ is cumulative distribution function of the standard normal distribution.

In light of (B.7), let us rephrase the UMP in terms of the ratio $\mathcal{P}_b/\mathcal{P}_{b+s}$. For convenience, let

$$W \equiv e^T = \frac{\mathcal{P}_b(\{\Upsilon\})}{\mathcal{P}_{b+s}(\{\Upsilon\})} . \quad (\text{B.12})$$

The critical region of the test is given by

$$W \leq W^* \equiv e^{T^*} . \quad (\text{B.13})$$

Furthermore, it follows from the asymptotic distribution of T , that the asymptotic distribution of W under $\mathbb{H}_{b+s}$ is Lognormal($\mu = -D_{\text{KL}}, \sigma^2$). Plugging these in (B.7), we can write the size of the UMP α^* , in the asymptotic limit as

$$\alpha^* \rightarrow (1 - \beta^*) \mathbb{E}[W \mid W \leq W^*] , \quad (\text{B.14})$$

for W distributed as-per Lognormal($\mu = -D_{\text{KL}}, \sigma^2$). The conditional expectation of W under the Lognormal distribution can be written as³¹

$$\mathbb{E}[W \mid W \leq W^*] = e^{\mu + (\sigma^2/2)} \frac{\Phi\left[\frac{\ln(W^*) - \mu - \sigma^2}{\sigma}\right]}{\Phi\left[\frac{\ln(W^*) - \mu}{\sigma}\right]} = e^{\mu + (\sigma^2/2)} \frac{\Phi\left[\frac{T^* - \mu - \sigma^2}{\sigma}\right]}{\Phi\left[\frac{T^* - \mu}{\sigma}\right]} . \quad (\text{B.15})$$

³⁰Since the number of events in the sample (and hence the number of terms in the summation) is random, here we invoke a special form of the central limit theorem for the sum of a random number of independent random variables [60, 61].

³¹This can be proved by making the substitution $t = (\ln(W) - (\mu + \sigma^2))/\sigma$ in the integral expression for the conditional expectation.

$\Phi \left[\frac{T^* - \mu}{\sigma} \right]$ asymptotes to $(1 - \beta^*)$, as per (B.11). On the other hand,

$$\Phi \left[\frac{T^* - \mu - \sigma^2}{\sigma} \right] \longrightarrow \Phi \left[\Phi^{-1}(1 - \beta^*) - \sigma \right] . \quad (\text{B.16})$$

In terms of the error function erf this can be written as

$$\Phi \left[\frac{T^* - \mu - \sigma^2}{\sigma} \right] \longrightarrow \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left[\operatorname{erf}^{-1}(1 - 2\beta^*) - \frac{\sigma}{\sqrt{2}} \right] . \quad (\text{B.17})$$

If $(1 - 2\beta^*) \leq (1 - \beta^*) < 1$, then $\operatorname{erf}^{-1}(1 - 2\beta^*)$ is finite (and fixed). Recall that σ^2 scales linearly with the integrated luminosity in the asymptotic limit. So, we can expand this expression up to the leading order in $1/\sigma$ for large σ , which gives

$$\Phi \left[\frac{T^* - \mu - \sigma^2}{\sigma} \right] \longrightarrow \frac{1}{\sigma\sqrt{2\pi}} \exp \left[- \left(\frac{\sigma}{\sqrt{2}} - \operatorname{erf}^{-1}(1 - 2\beta^*) \right)^2 \right] \quad (\text{B.18})$$

$$= e^{-\sigma^2/2} \frac{1}{\sigma\sqrt{2\pi}} \exp \left[\sqrt{2}\sigma \operatorname{erf}^{-1}(1 - 2\beta^*) - \left[\operatorname{erf}^{-1}(1 - 2\beta^*) \right]^2 \right] . \quad (\text{B.19})$$

Plugging (B.11), (B.15), and (B.19) into (B.14), and using $\mu = -D_{\text{KL}}$ we get

$$\alpha^* \longrightarrow e^{-D_{\text{KL}}} \frac{1}{\sigma\sqrt{2\pi}} \exp \left[\sqrt{2}\sigma \operatorname{erf}^{-1}(1 - 2\beta^*) - \left[\operatorname{erf}^{-1}(1 - 2\beta^*) \right]^2 \right] \quad (\text{B.20})$$

Noting that a) $\operatorname{erf}^{-1}(1 - 2\beta^*)$ is finite and fixed, b) D_{KL} scales linearly with the integrated luminosity, and c) σ scales as square-root of the integrated luminosity, we can write

$$\alpha^* \longrightarrow e^{-D_{\text{KL}} + o(D_{\text{KL}})} . \quad (\text{B.21})$$

which completes the proof of (B.1).

Statement (B.2) can be proved in a similar fashion. Alternatively, it can also be derived from (B.1) by making the substitutions $\alpha^* \rightarrow \beta^*$ (Type I error $\rightarrow$ Type II error) and $D_{\text{KL}} \rightarrow D_{\text{revKL}}$.

C Proof of correctness of the iterative training algorithm

In this section we will prove the correctness as well as convergence of our event categorizer training prescription. The key ingredients to the proof are the following two properties

- (a) $\Lambda_{\text{stat}}(U, V)$ is linear in V .
- (b) $\Lambda_{\text{stat}}(U, V)$ has a global maximum in U at $U = V$.

The linearity property (a) is by construction of Λ_{stat} -s from D_{stat} -s. The global maximum property (b) is from the nature of D_{stat} -s, which penalize mixing of regions of phase-space with different $p(\mathbf{e})$ values.

In each iterative step, from the current categorizer, η , a new categorizer, $\bar{\eta}$, is constructed. Let $p_c(\mathbf{x})$ and $\bar{p}_c(\mathbf{x})$ be the signal-fraction functions of η and $\bar{\eta}$ respectively. Recall that $\bar{\eta}$ rejects an event (if allowed) if $\Lambda_{\text{stat}}(p_c(\mathbf{x}), p(\mathbf{e})) < 0$ for all categories c . Otherwise, $\bar{\eta}$ places each event in the category that maximizes $\Lambda_{\text{stat}}(p_c(\mathbf{x}), p(\mathbf{e}))$. From the construction of $\bar{\eta}$ it is clear that

$$\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \eta(\mathbf{e})) \Lambda_{\text{stat}}(p_c(\mathbf{x}), p(\mathbf{e})) \right] \leq \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(p_c(\mathbf{x}), p(\mathbf{e})) \right] \quad (\text{C.1})$$

By treating “averaging over all events” as averaging over events with a given value of $(c, \mathbf{x})$ followed by an (appropriately weighted) average over $(c, \mathbf{x})$ values, we can write

$$\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(p_c(\mathbf{x}), p(\mathbf{e})) \right] = \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(p_c(\mathbf{x}), \bar{p}_c(\mathbf{x})) \right] \quad (\text{C.2})$$

Here we have used the fact that all Λ_{stat} -s are linear in the second variable. Though we don’t explicitly derive eq. (C.2) here, it can be shown to follow from eq. (2.13). Next, using the fact that $\Lambda_{\text{stat}}(A, B) \leq \Lambda_{\text{stat}}(B, B)$ for all $A, B \in [0, 1)$, we can write

$$\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(p_c(\mathbf{x}), \bar{p}_c(\mathbf{x})) \right] \leq \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(\bar{p}_c(\mathbf{x}), \bar{p}_c(\mathbf{x})) \right] \quad (\text{C.3})$$

Using eq. (2.13) and the linearity of Λ_{stat} in the second variable again, we can bring $p(\mathbf{e})$ back into the expression as

$$\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(\bar{p}_c(\mathbf{x}), \bar{p}_c(\mathbf{x})) \right] = \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(\bar{p}_c(\mathbf{x}), p(\mathbf{e})) \right] \quad (\text{C.4})$$

Combining eqns (C.1)–(C.4), we can see that

$$\text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \eta(\mathbf{e})) \Lambda_{\text{stat}}(p_c(\mathbf{x}), p(\mathbf{e})) \right] \leq \text{MEAN}_{\mathbf{e} \sim \mathbb{H}_{b+s}} \left[\sum_{c=1}^C \delta(c, \bar{\eta}(\mathbf{e})) \Lambda_{\text{stat}}(\bar{p}_c(\mathbf{x}), p(\mathbf{e})) \right] \quad (\text{C.5})$$

This can be rewritten using eq. (2.20) as

$$\left\{ \tilde{D}_{\text{Neym}\chi^2} \text{ corresponding to } \eta \right\} \leq \left\{ \tilde{D}_{\text{Neym}\chi^2} \text{ corresponding to } \bar{\eta} \right\} . \quad (\text{C.6})$$

Eq. (C.6) shows that the categorizer “improves” in each iterative step. The only caveat here is the assumption that $p_c(\mathbf{x})$ can be estimated accurately enough at each step.

To see that the algorithm will converge in a finite number of steps, note that during the training process a categorizer in the current step is defined completely by the assignment of training events to categories in the previous step. If the categorizer has to improve in each step, it has to be different in each step (no cycles), and there is only a finite number of possible category assignments for training events. So, in a finite number of iterations, the algorithm should converge on a categorizer that maximizes (the estimate of) $\tilde{D}_{\text{stat}}$ *locally*.

A caveat to this argument is that due to artifacts in the unbinned estimation of the signal-fraction functions $p_c(\mathbf{x})$, it is possible to run into cycles near the end of the training process. So it may make sense to terminate the iteration process when the fraction of events reassigned in some step drops below a certain threshold or if the (absolute or relative) gain in the estimated $\tilde{D}_{\text{stat}}$ drops below a certain threshold. For a binned analysis the convergence argument presented above holds without caveats, and it is guaranteed that $\bar{\eta}$ will be the same as η in some training step — although it may take an unfeasible number of iterations to reach this stage.

The time complexity of the training process is $\mathcal{O}(CNi + f(C, \mathcal{N})i)$ where $\mathcal{N}$ is the number of training events, C is the number of categories, i is the number of iterative steps. $f(C, \mathcal{N})$ is the cost associated with estimating $p_c(\mathbf{x})$, for each training event and c , from the training data (line 6 and line 8 (called from line 5) of algorithm 1). For binned analyses with M $\mathbf{x}$ -bins, $f(\mathcal{N}, C)$ is $\mathcal{O}(CM + CN)$, so the overall time complexity for training is $\mathcal{O}(CNi)$ (assuming $\mathcal{N} > M$, i.e., there will be more training events than bins).

For unbinned analyses, it is possible to estimate $p_c(\mathbf{x})$ using kernel regression (with appropriate choice of kernels) in $\mathcal{O}(2^d CN \log \mathcal{N})$ time where d is the dimensionality of $\mathbf{x}$ (using a k-d tree based implementation). This leads to an overall training time complexity of $\mathcal{O}(2^d CN \log \mathcal{N} i)$.

While the number of iterations i needed for convergence in the worst case is superpolynomial [63] (unbounded by any polynomial), in practice, for reasonable values of C and d , the value of $\tilde{D}_{\text{stat}}$ flattens out after a few tens of iterations — just like with Lloyd’s k-means clustering algorithm. So, with appropriate termination criteria, the iterative training process can be implemented to practically have a time complexity of $\mathcal{O}(CN)$ when using binned and $\mathcal{O}(CN \log \mathcal{N})$ when using unbinned estimation of p_c (ignoring i and 2^d).

D A comparison of the different distance measures

In a given iterative step of training, each event $\mathbf{e}$ will be categorized into one of the categories based on their $p_c(\mathbf{x}(\mathbf{e}))$ values from the previous step — event $\mathbf{e}$ will be placed into the one of the categories with previous $p_c(\mathbf{x}(\mathbf{e}))$ closest to $p(\mathbf{e})$ on either side (if event rejection is allowed, rejected events can be thought of as forming a category of their own with $p_c(\mathbf{x}) = 0$). The only difference between using different measures is in the threshold value of $p(\mathbf{e})$, for a pair of categories, above which events will be placed in the “higher purity” category.

Let c_1 and c_2 be two adjacent categories when ordered by $p_c(\mathbf{x})$ for a given $\mathbf{x}$, with $p_{c_1}(\mathbf{x}) < p_{c_2}(\mathbf{x})$. At the threshold value of $p(\mathbf{e})$, say $p_{\text{stat}}^{\text{thresh}}$, between $p_{c_1}(\mathbf{x})$ and $p_{c_2}(\mathbf{x})$, $\Lambda_{\text{stat}}[p_{c_1}(\mathbf{x}), p_{\text{stat}}^{\text{thresh}}] = \Lambda_{\text{stat}}[p_{c_2}(\mathbf{x}), p_{\text{stat}}^{\text{thresh}}]$. For values of $p(\mathbf{e})$ greater (lesser) than $p_{\text{stat}}^{\text{thresh}}$, $\Lambda_{\text{stat}}(p_{c_2}(\mathbf{x}), p(\mathbf{e}))$ is greater (lesser) than $\Lambda_{\text{stat}}(p_{c_1}(\mathbf{x}), p(\mathbf{e}))$.

Formally, let us define the threshold value $p_{\text{stat}}^{\text{thresh}}(p_1, p_2)$ of $p(\mathbf{e})$, at a given value of $\mathbf{x}$, between two categories with $p_{c_1}(\mathbf{x}) = p_1$ and $p_{c_2}(\mathbf{x}) = p_2$, as satisfying

$$\min(p_1, p_2) \leq p_{\text{stat}}^{\text{thresh}}(p_1, p_2) \leq \max(p_1, p_2) , \quad (\text{D.1a})$$

$$\Lambda_{\text{stat}}(p_1, p_{\text{stat}}^{\text{thresh}}(p_1, p_2)) = \Lambda_{\text{stat}}(p_2, p_{\text{stat}}^{\text{thresh}}(p_1, p_2)) . \quad (\text{D.1b})$$

$p_{\text{stat}}^{\text{thresh}}$ is related to $p_{\text{stat}}^{\text{sel_thresh}}$ defined eq. (2.23) as $p_{\text{stat}}^{\text{sel_thresh}}(p_c) = p_{\text{stat}}^{\text{thresh}}(0, p_c)$.

Figure 14. Plots of $p_{\text{stat}}^{\text{thresh}}(p_1, p_2)$ for the different statistical distances (middle column) arranged in decreasing order of strictness (top to bottom). The right column shows the ratio between $p_{\text{stat}}^{\text{thresh}}(p_1, p_2)$ for consecutive distances, in decreasing order of strictness. The ratios being ≥ 1 demonstrates the correctness of their ordering.

It can be proved that $\forall (p_1, p_2) \in [0, 1) \times [0, 1)$,

$$p_{\text{Neym}\chi^2}^{\text{thresh}} \leq p_{\text{revKL}}^{\text{thresh}} \leq p_{\text{B}}^{\text{thresh}} \leq p_{\text{J}}^{\text{thresh}} \leq p_{\text{KL}}^{\text{thresh}} \leq p_{\text{Pear}\chi^2}^{\text{thresh}} \quad , \quad (\text{D.2})$$

with equality iff $p_1 = p_2$.

This property is written in eq. (3.3) as the “strictness” order of the statistical distances, with $\tilde{D}_{\text{Neym}\chi^2}$ and $\tilde{D}_{\text{Pear}\chi^2}$ being the least and most strict respectively. In figure 14 we plot $p_{\text{stat}}^{\text{thresh}}(p_1, p_2)$ for the different statistical distances (middle column) arranged in decreasing order of strictness (top to bottom). Figure 14 also shows the ratio between $p_{\text{stat}}^{\text{thresh}}(p_1, p_2)$ for consecutive distances, in decreasing order of strictness (right panel). The ratios being greater than or equal to one demonstrates the correctness of the ordering in eq. (D.2).

Finally we want to point out that there is nothing sacred about the six distance measures used in this paper. Any measure of distance that is joint-convex in the distributions $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$, and proportional to the Fisher-Rao metric distance between them in the $\mathbb{H}_{b+s} \rightarrow \mathbb{H}_b$ limit, can be used for training. The six distances provided in this paper are just a few named distances known in the statistics literature for which we were able to find usable closed form expressions in terms of the number density functions $s_c(\mathbf{x})$ and $b_c(\mathbf{x})$. For example, any linear combination of $\tilde{D}_{\text{stat}}$ ’s from eq. (2.19) can also be used to train categorizers, along with the same linear combination of eventwise contributions Λ_{stat} from eq. (2.21), although this much choice is probably overkill for most applications.

On a related note, we would like to point to the following approximation for use in estimating the statistical significance of deviations from expectation:

$$\begin{aligned} & \left[\sum_i 2 \left[-(\text{obs}_i - \text{exp}_i) + \text{obs}_i \ln \left(\frac{\text{obs}_i}{\text{exp}_i} \right) \right] \right] \\ & \approx \left[\frac{1}{3} \sum_i \frac{(\text{obs}_i - \text{exp}_i)^2}{\text{obs}_i} + \frac{2}{3} \sum_i \frac{(\text{obs}_i - \text{exp}_i)^2}{\text{exp}_i} \right] \end{aligned} \quad (\text{D.3})$$

Here the linear combination of Neyman- χ^2 and Pearson- χ^2 test statistics on the right-hand side, dubbed the “combined Neyman-Pearson χ^2 ” in [64], approximates the left-hand side better than they do individually. In terms of our statistical distances between $\mathbb{H}_b$ and $\mathbb{H}_{b+s}$, this approximation can be written as $\tilde{D}_{\text{KL}} \approx (\tilde{D}_{\text{Neym}\chi^2} + 2 \tilde{D}_{\text{Pear}\chi^2}) / 3$.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] D. Guest, K. Cranmer and D. Whiteson, *Deep learning and its application to LHC physics*, *Ann. Rev. Nucl. Part. Sci.* **68** (2018) 161 [[arXiv:1806.11484](https://arxiv.org/abs/1806.11484)] [[INSPIRE](#)].
- [2] K. Albertsson et al., *Machine learning in high energy physics community white paper*, *J. Phys. Conf. Ser.* **1085** (2018) 022008 [[arXiv:1807.02876](https://arxiv.org/abs/1807.02876)] [[INSPIRE](#)].
- [3] D. Bourilkov, *Machine and deep learning applications in particle physics*, to appear.

- [4] CMS collaboration, *Flexible trigger menu implementation on the global trigger for the CMS level-1 trigger upgrade*, *J. Phys. Conf. Ser.* **898** (2017) 032033 [[INSPIRE](#)].
- [5] ATLAS collaboration, *The ATLAS trigger menu design for higher luminosities in Run 2*, *EPJ Web Conf.* **214** (2019) 01031 [[INSPIRE](#)].
- [6] F. Rosenblatt, *The perceptron: a probabilistic model for information storage and organization in the brain*, *Psychol. Rev.* **65** (1958) 386.
- [7] L. Breiman, J.H. Friedman, R.A. Olshen and C.J. Stone, *Classification and regression trees*, Wadsworth and Brooks, Monterey U.S.A. (1984).
- [8] Y. Freund and R.E. Schapire, *Experiments with a new boosting algorithm*, in the proceedings of the 13th *International Conference on Machine Learning*, July 3–6, Bari, Italy (1996).
- [9] K.T. Matchev and P. Shyamsundar, *ThickBrick: optimal event selection and categorization in high energy physics. Part II. Parameter measurement*, in preparation.
- [10] K.T. Matchev and P. Shyamsundar, *ThickBrick: optimal event selection and categorization in high energy physics. Part III. Nuisance parameters and systematic uncertainties*, in preparation.
- [11] K. Kondo, *Dynamical likelihood method for reconstruction of events with missing momentum. 1: method and toy models*, *J. Phys. Soc. Jap.* **57** (1988) 4126 [[INSPIRE](#)].
- [12] R.H. Dalitz and G.R. Goldstein, *The decay and polarization properties of the top quark*, *Phys. Rev. D* **45** (1992) 1531
- [13] J.S. Gainer, J. Lykken, K.T. Matchev, S. Mrenna and M. Park, *The matrix element method: past, present, and future*, [arXiv:1307.3546](#) [[INSPIRE](#)].
- [14] K. Cranmer, J. Pavez and G. Louppe, *Approximating likelihood ratios with calibrated discriminative classifiers*, [arXiv:1506.02169](#) [[INSPIRE](#)].
- [15] G. Louppe, K. Cranmer and J. Pavez, *carl: a likelihood-free inference toolbox*, *J. Open Source Softw.* **1** (2016) 11.
- [16] J. Brehmer, K. Cranmer, I. Espejo, F. Kling, G. Louppe and J. Pavez, *Effective LHC measurements with matrix elements and machine learning*, *J. Phys. Conf. Ser.* **1525** (2020) 012022 [[arXiv:1906.01578](#)] [[INSPIRE](#)].
- [17] J. Brehmer, F. Kling, I. Espejo and K. Cranmer, *MadMiner: machine learning-based inference for particle physics*, *Comput. Softw. Big Sci.* **4** (2020) 3 [[arXiv:1907.10621](#)] [[INSPIRE](#)].
- [18] S. Chen, A. Glioti, G. Panico and A. Wulzer, *Parametrized classifiers for optimal EFT sensitivity*, [arXiv:2007.10356](#) [[INSPIRE](#)].
- [19] T.M. Cover and J. A. Thomas, *Elements of information theory*, 2nd edition, Wiley, U.S.A. (2006).
- [20] J. Neyman and E.S. Pearson, *On the problem of the most efficient tests of statistical hypotheses*, *Phil. Trans. Roy Soc. London A* **231** (1933) 289.
- [21] J.L.W.V. Jensen, *Sur les fonctions convexes et les inégalités entre les valeurs moyennes*, *Acta Math.* **30** (1906) 175.
- [22] S. Whiteson and D. Whiteson, *Stochastic optimization for collision selection in high energy physics*, in the proceedings of the 19th *National Conference on Innovative Applications of Artificial Intelligence (IAAI-07)*, July 22–26, Vancouver, Canada (2007).

- [23] S. Whiteson and D. Whiteson, *Machine learning for event selection in high energy physics*, *Eng. Appl. Artif. Intell.* **22** (2009) 1203.
- [24] K.O. Stanley and R. Miikkulainen, *Evolving neural networks through augmenting topologies*, *Evolut. Comput.* **10** (2002) 99.
- [25] A. Elwood and D. Krücker, *Direct optimisation of the discovery significance when training neural networks to search for new physics in particle colliders*, [arXiv:1806.00322](#) [[INSPIRE](#)].
- [26] CMS collaboration, *Machine learning techniques in the CMS search for Higgs decays to dimuons*, *EPJ Web Conf.* **214** (2019) 06002 [[INSPIRE](#)].
- [27] A. Valassi, *Binary classifier metrics for optimizing HEP event selection*, *EPJ Web Conf.* **214** (2019) 06004 [[INSPIRE](#)].
- [28] A. Valassi, *Optimising HEP parameter fits via Monte Carlo weight derivative regression*, *EPJ Web Conf.* **245** (2020) 06038 [[arXiv:2003.12853](#)] [[INSPIRE](#)].
- [29] A. Alves, T. Ghosh and K. Sinha, *Can we discover double Higgs production at the LHC?*, *Phys. Rev. D* **96** (2017) 035022 [[arXiv:1704.07395](#)] [[INSPIRE](#)].
- [30] A. Alves, T. Ghosh, H.-K. Guo and K. Sinha, *Resonant di-Higgs production at gravitational wave benchmarks: a collider study using machine learning*, *JHEP* **12** (2018) 070 [[arXiv:1808.08974](#)] [[INSPIRE](#)].
- [31] J. Dolen, P. Harris, S. Marzani, S. Rappoccio and N. Tran, *Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure*, *JHEP* **05** (2016) 156 [[arXiv:1603.00027](#)] [[INSPIRE](#)].
- [32] J.A. Aguilar-Saavedra, J.H. Collins and R.K. Mishra, *A generic anti-QCD jet tagger*, *JHEP* **11** (2017) 163 [[arXiv:1709.01087](#)] [[INSPIRE](#)].
- [33] S. Chang, T. Cohen and B. Ostdiek, *What is the machine learning?*, *Phys. Rev. D* **97** (2018) 056009 [[arXiv:1709.10106](#)] [[INSPIRE](#)].
- [34] J. Stevens and M. Williams, *uBoost: a boosting method for producing uniform selection efficiencies from multivariate classifiers*, *2013 JINST* **8** P12013 [[arXiv:1305.7248](#)] [[INSPIRE](#)].
- [35] C. Shimmin et al., *Decorrelated jet substructure tagging using adversarial neural networks*, *Phys. Rev. D* **96** (2017) 074034 [[arXiv:1703.03507](#)] [[INSPIRE](#)].
- [36] L. Bradshaw, R.K. Mishra, A. Mitridate and B. Ostdiek, *Mass Agnostic Jet Taggers*, *SciPost Phys.* **8** (2020) 011 [[arXiv:1908.08959](#)] [[INSPIRE](#)].
- [37] M. Pivk and F.R. Le Diberder, *SPlot: a statistical tool to unfold data distributions*, *Nucl. Instrum. Meth. A* **555** (2005) 356 [[physics/0402083](#)] [[INSPIRE](#)].
- [38] ThickBrick package, <https://prasanthcakewalk.gitlab.io/thickbrick/>.
- [39] J. Neyman, *Contribution to the theory of chi-square test*, in the proceedings of the *First Berkeley Symposium on Mathematical Statistics and Probability*, August 13–18, Berkeley U.S.A. (1945).
- [40] K. Pearson, *On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling*, *Philos. Mag.* **50** (1900) 157.
- [41] S. Kullback and R.A. Leibler, *On information and sufficiency*, *Ann. Math. Stat.* **22** (1951) 79.

- [42] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554 [Erratum *ibid.* **73** (2013) 2501] [[arXiv:1007.1727](#)] [[INSPIRE](#)].
- [43] H. Jeffreys, *An invariant form for the prior probability in estimation problems*, *Proc. Roy Soc. London A* **186** (1946) 453.
- [44] T. Kailath, *The divergence and Bhattacharyya distance measures in signal selection*, *IEEE Trans. Comm. Techn.* **15** (1967) 52.
- [45] S.P. Lloyd, *Least squares quantization in PCM*, *IEEE Trans. Inf. Theor.* **28** (1982) 129.
- [46] E. Forgy, *Cluster analysis of multivariate data: efficiency versus interpretability of classification*, *Biometrics* **21** (1965) 768.
- [47] N. Wardle, *Per-event significance indicator to visualise significant events*, [arXiv:1902.08508](#) [[INSPIRE](#)].
- [48] J.S. Gainer, J. Lykken, K.T. Matchev, S. Mrenna and M. Park, *Exploring theory space with Monte Carlo reweighting*, *JHEP* **10** (2014) 078 [[arXiv:1404.7129](#)] [[INSPIRE](#)].
- [49] O. Mattelaer, *On the maximal use of Monte Carlo samples: re-weighting events at NLO accuracy*, *Eur. Phys. J. C* **76** (2016) 674 [[arXiv:1607.00763](#)] [[INSPIRE](#)].
- [50] L. Bottou and Y. Bengio, *Convergence properties of the k-means algorithms*, in the proceedings of the 7th International Conference on Neural Information Processing Systems (*NIPS'94*), Denver, U.S.A. (1994).
- [51] D. Sculley, *Web-scale k-means clustering*, in the proceedings of the 19th International Conference on the World Wide Web, April 26–30, Raleigh, U.S.A. (2010).
- [52] K.T. Matchev, A. Roman and P. Shyamsundar, *ThickBrick tutorial for PyHEP 2020*, <https://github.com/prasanthcakewalk/ThickBrick-Tutorial-PyHEP-2020>.
- [53] K.T. Matchev, A. Roman and P. Shyamsundar, *Optimizing physics searches at the LHC using ThickBrick*, in preparation.
- [54] E.A. Nadaraya, *On estimating regression*, *Theory Prob. Appl.* **9** (1964) 141.
- [55] G.S. Watson, *Smooth regression analysis*, *Sankhyā Ser. A* **26** (1964) 359.
- [56] G. Louppe, M. Kagan and K. Cranmer, *Learning to pivot with adversarial networks*, [arXiv:1611.01046](#) [[INSPIRE](#)].
- [57] L.M. Dery, B. Nachman, F. Rubbo and A. Schwartzman, *Weakly supervised classification in high energy physics*, *JHEP* **05** (2017) 145 [[arXiv:1702.00414](#)] [[INSPIRE](#)].
- [58] E.M. Metodiev, B. Nachman and J. Thaler, *Classification without labels: learning from mixed samples in high energy physics*, *JHEP* **10** (2017) 174 [[arXiv:1708.02949](#)] [[INSPIRE](#)].
- [59] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, *Learning to classify from impure samples with high-dimensional data*, *Phys. Rev. D* **98** (2018) 011502 [[arXiv:1801.10158](#)] [[INSPIRE](#)].
- [60] A. Rényi, *On the central limit theorem for the sum of a random number of independent random variables*, *Acta Math. Acad. Scient. Hung.* **11** (1963) 97.
- [61] J.R. Blum, D.L. Hanson and J.I. Rosenblatt, *On the central limit theorem for the sum of a random number of independent random variables*, *Z. Wahrscheinlichkeitstheorie verw. Gebiete* **1** (1963) 389.

- [62] C.M. Jones, A.A. Zhigljavsky, *Approximating the negative moments of the Poisson distribution*, *Stat. Probab. Lett.* **66** (2004) 171.
- [63] D. Arthur and S. Vassilvitskii, *How slow is the k-means method?*, in the proceedings of the 22nd *Annual Symposium on Computational Geometry (SoCG06)*, June 5–7, Sedona, U.S.A. (2006).
- [64] X. Ji, W. Gu, X. Qian, H. Wei and C. Zhang, *Combined Neyman-Pearson χ -square: An improved approximation to the Poisson-likelihood χ -square*, *Nucl. Instrum. Meth. A* **961** (2020) 163677 [[arXiv:1903.07185](#)] [[INSPIRE](#)].