
Toward Model-Agnostic Detection of New Physics Using Data-Driven Signal Regions

Soheun Yi

Department of Statistics and Data Science
Carnegie Mellon University
Pittsburgh, PA 15213
soheuny@andrew.cmu.edu

John Alison

Department of Physics
Carnegie Mellon University
Pittsburgh, PA 15213
johnalison@cmu.edu

Mikael Kuusela

Department of Statistics and Data Science
Carnegie Mellon University
Pittsburgh, PA 15213
mkuusela@andrew.cmu.edu

Abstract

In the search for new particles in high-energy physics, it is crucial to select the Signal Region (SR) in such a way that it is enriched with signal events if they are present. While most existing search methods set the region relying on prior domain knowledge, it may be unavailable for a completely novel particle that falls outside the current scope of understanding. We address this issue by proposing a method built upon a model-agnostic but often realistic assumption about the localized topology of the signal events, in which they are concentrated in a certain area of the feature space. Considering the signal component as a localized high-frequency feature, our approach employs the notion of a low-pass filter. We define the SR as an area which is most affected when the observed events are smeared with additive random noise. We overcome challenges in density estimation in the high-dimensional feature space by learning the density ratio of events that potentially include a signal to the complementary observation of events that closely resemble the target events but are free of any signals. By applying our method to simulated $HH \rightarrow 4b$ events, we demonstrate that the method can efficiently identify a data-driven SR in a high-dimensional feature space in which a high portion of signal events concentrate.

1 Introduction

High-energy physics is a subfield of physics that studies fundamental particles and their interactions to understand how matter and radiation in the universe are formed. At its core, the Standard Model (SM) plays a central role in explaining the fundamental particles and their interactions. The remarkable discovery of the Higgs boson by collaborations at the European Organization for Nuclear Research (CERN) [1, 6] and its inclusion in the SM leads to a natural question about the possibility of new particles and interactions beyond the SM. An important challenge in searching for new physics events is to estimate the background distribution expected from the SM, which is to be compared with the observed data to identify the existence of new signals. Perhaps the most prominent approaches to background estimation are ABCD and sideband methods, and their variants [1, 3, 6, 8]. At the center of these approaches, one selects the *Signal Region* (SR) where signal events are expected to be concentrated, and the *Control Region* (CR) where signal events are expected to be nearly absent.




Figure 1: Depiction of the variables used to represent a particle jet. p_T is the transverse momentum, ϕ is the azimuthal angle, and θ is the polar angle. The pseudorapidity is given by $\eta = -\log(\tan(\theta/2))$.

Then, the background distribution is estimated from the CR and extrapolated to the SR to compare with the observed data. Most existing methods invoke knowledge about the expected signals to select the SR and CR. In detecting the Higgs boson, for example, the SR is configured to include the mass range where the Higgs boson is expected to appear and the CR is configured to exclude the SR [1, 3, 6].

For a completely new type of particle, however, we are likely to lack specific prior knowledge about the signal events, which has led to extensive recent work on methods for model-agnostic signal searches [4, 9, 11–14, 16]. In this case, setting up the SR and CR based on existing knowledge about the signal is not feasible. This motivates us to develop a method to define the SR and CR only with the observed data without relying heavily on prior domain knowledge about the signals. Rather, we observe that signals tend to be localized in the feature space, which is an assumption that has been commonly exploited in certain methods, namely “bump hunting” [2, 5, 9, 10].

For instance, the Higgs boson is expected to appear in a particular invariant mass range around 125 GeV. Motivated by this observation, we make a minimal assumption that the signal events are localized in the (potentially high-dimensional) feature space and use this to propose a data-driven way of selecting the SR and CR.

2 Problem Formulation

When two protons collide, they form highly unstable particles that decay within an extremely short time frame, making them impossible to detect directly. Some of these transient particles decay into quarks, which interact with the vacuum to produce collimated sprays of stable particles, referred to as “jets”. Among all possible final states, we focus on events resulting in four jets.

Being able to detect jets efficiently, the Compact Muon Solenoid (CMS) detector at CERN records the momentum and the invariant mass (or energy) of the particles produced in proton-proton collisions. For the momentum of an observed particle, let (r, θ, ϕ) be its representation in spherical coordinates, where $\theta \in [0, \pi)$ is the polar angle and $\phi \in [0, 2\pi)$ is the azimuthal angle. As visually explained in Figure 1, the momentum vector is represented by a three-dimensional vector (p_T, η, ϕ) ; $p_T := r \sin \theta$ is the *transverse momentum*, or the magnitude of the momentum projected onto the xy -plane, $\eta := -\log \tan(\theta/2)$ is the *pseudorapidity*. Writing the rest mass of a particle as m , the measured features of an observed particle are represented as a four-vector $(p_T, \eta, \phi, m) \in \mathbb{R}^4$. As we are concerned with collisions resulting in 4 jets, the corresponding events can be represented with $x = (p_T^i, \eta^i, \phi^i, m^i)_{i=1}^4 \in \mathbb{R}^{16}$, forming a 16-dimensional feature space $\mathcal{X} \subset \mathbb{R}^{16}$.

The SM predicts a rare chance of a proton-proton collision producing two Higgs bosons, and in such case, they are most likely to decay into four b -quarks which produce specific types of jets called “ b -jets”. This instance of interactions of particles is abbreviated as $HH \rightarrow 4b$ as it incorporates the process of two Higgs bosons decaying into four b -jets. We refer to proton-proton collisions that are tagged to have produced four b -jets as $4b$ events, and among them, we are interested in identifying the existence of events corresponding to $HH \rightarrow 4b$. In the sequel, we refer to these events as the *signal events*, and the rest of the $4b$ events as the *background $4b$ events*. In addition to the $4b$ events, we have

observations of events that are tagged to have produced three b -jets; we refer to these events as $3b$ events and make the assumption that the signal we are looking for is not present among the $3b$ events. We leverage the density ratio of $4b$ events to $3b$ events to select the SR and CR in the 16-dimensional feature space. We elaborate further on this in Section 3.

We have access to n samples of $3b$ and m samples of $4b$ events, denoted by $\{X_{3b}^1, \dots, X_{3b}^n\}$ and $\{X_{4b}^1, \dots, X_{4b}^m\}$, respectively. Denote these distributions by

$$X_{3b}^1, \dots, X_{3b}^n \stackrel{\text{iid}}{\sim} P_{3b}, \quad X_{4b}^1, \dots, X_{4b}^m \stackrel{\text{iid}}{\sim} P_{4b} = (1 - \epsilon)B_{4b} + \epsilon S, \quad (1)$$

where B_{4b} and S are distributions (or probability measures) of background $4b$ and signal events, respectively. We call ϵ the *signal ratio*, which is the fraction of signal events among the observed events. Since a nonzero ϵ indicates the existence of a signal in the data, the problem of discovering new physics boils down to testing the following hypotheses:

$$H_0: \epsilon = 0 \quad \text{vs.} \quad H_1: \epsilon > 0. \quad (2)$$

Under the existence of the signal process, our goal is to select the SR, denoted by $\mathcal{X}_s \subset \mathcal{X}$, to be rich in signals relative to its size, which can be expressed as amplifying $S(\mathcal{X}_s)$ compared to $P_{4b}(\mathcal{X}_s)$. This is important for improving the power of the test conducted on the signal region since a higher concentration of signal should lead to a greater deviation from the null hypothesis.

3 Methodology

Learning Event Representations. An important factor to consider for our collision events is how the four jets are naturally paired into two sets of two jets, according to the intermediate particle from which each jet originates. In addition, there are symmetries in the feature space (reflection symmetries in ϕ and η , as well as a rotational symmetry in ϕ) to be considered. To address the complicated topology in the original feature space posed by these aspects, we exploit the classifier architecture proposed by Manole et al. [15], which is designed to handle both the pairing of the jets and the symmetries in the feature space. Informally speaking, the classifier is trained to separate $3b$ events from $4b$ events and the last layer of the classifier calculates the probability of an event being a $4b$ event as opposed to a $3b$ event from inputs calculated by the previous layers. We use these inputs to that last layer as the representations of events with the expectation that they exhibit a simpler topology. We denote the representation of an event x as $\zeta(x)$, or simply ζ when the context is clear.

Finding the Data-Driven Signal Region. As mentioned in Section 1, we leverage the assumption that the signal events are localized in the feature space, or the space of representations ζ . Recall that signal events of our interest are tagged as $4b$ events. Hence, an area where signals are concentrated would display a higher density of $4b$ events compared to $3b$ events. In the idealized case of the density ratio of $4b$ events to $3b$ events being uniform across the feature space except at those locations, simply selecting areas near the peak of (or area with the highest value of) the density ratio would be sufficient to identify the SR.

Unfortunately, this is not the case in reality. It is possible that the density ratio in some areas of the feature space is higher than the ratio in areas where the signal is concentrated, making those regions merely local maxima of the density ratio. A plot on the upper left of Figure 2 illustrates this issue in a low-dimensional toy example; the center of the signal does not necessarily have the highest value of the density ratio. To address this issue, we leverage the fact that $3b$ events possess similar (but not identical) kinematics to that of background $4b$ events [7], which allows us to postulate that the density ratio of background $4b$ events to $3b$ events should only consist of “low-frequency” features. Therefore, if the density ratio of *all* $4b$ events to $3b$ events possesses a localized “high-frequency” feature, it would suggest the presence of a signal in the feature space. In other words, the density ratio of all $4b$ events to $3b$ events would show a local maximum (or a peak) in the feature space where signal events are concentrated.

We exploit this topology to find the SR and CR in the high-dimensional feature space. Motivated by the notion of a low-pass filter, we consider extracting low-frequency components of the $4b$ and $3b$ densities and then comparing the resulting density ratio with the original one to search for high-frequency components. If we take a convolution kernel (e.g., Gaussian kernel) with adequate bandwidth and convolve each density with it, we can obtain a smoothed version of the density ratio




Figure 2: $P_{3b} = \mathcal{N}(1, 4^2)$, $P_{4b} = 0.95B_{4b} + 0.05S = 0.95\mathcal{N}(-1, 4^2) + 0.05\mathcal{N}(7, 0.5^2)$, and $K = \mathcal{N}(0, 2^2)$. Signal events (with distribution S) are concentrated around $\zeta = 7$ (dashed lines). Looking at the largest values of γ (upper right) does not identify the center of signal events, but the ratio $\gamma/\tilde{\gamma}$ (lower right) does.

that reflects only low-frequency components in the densities with features corresponding to the signal suppressed. In mathematical terms, we define the density ratio and the smoothed density ratio as

$$\gamma(\zeta) := \frac{p_{4b}(\zeta)}{p_{3b}(\zeta)}, \quad \tilde{\gamma}(\zeta) := \frac{(p_{4b} * K)(\zeta)}{(p_{3b} * K)(\zeta)},$$

where p_{3b} (resp. p_{4b}) is the density of $3b$ (resp. $4b$) events in the representation space, $*$ denotes the convolution operation, and K is a convolution kernel. Remarkably, $\tilde{\gamma}$ can be efficiently estimated by learning a classifier without directly evaluating the convolution operation. In particular, we can estimate it by training a classifier to distinguish $(Z_{3b} + \mathcal{E}, 0)$ and $(Z_{4b} + \mathcal{E}, 1)$, where Z_{3b} and Z_{4b} are the representations of $3b$ and $4b$ events, respectively, and $\mathcal{E} \sim K$ is random noise.

A toy example shown in Figure 2 illustrates the concept of the smoothed density ratio. One can observe that the center of signal events cannot be identified by selecting a region with the highest value of γ , but rather by considering the ratio $\gamma/\tilde{\gamma}$. To this end, we define the data-driven SR as

$$\mathcal{X}_s := \{x \mid \gamma(\zeta(x))/\tilde{\gamma}(\zeta(x)) \geq \tau_s\}. \quad (3)$$

The threshold τ_s is specifically designed to control the size of the SR as measured by the proportion of $4b$ events within those regions.

4 Results

In training classifiers to estimate density ratios, we used $3b$ event samples of size $n \in \{10^5, 10^6\}$ and the same size of $4b$ events. We used 75% and 6.25% of all the samples to estimate γ and $\tilde{\gamma}$, respectively¹. To generate the training dataset for the smoothed density ratio $\tilde{\gamma}$, Gaussian noise was added to the representations of each event, with noise scale in each dimension set to $\eta \in \{0.01, 0.1, 1\}$ times the length of the range of the corresponding representation.

Recall that our goal is to enrich the SR (i.e., \mathcal{X}_s) with the signal, that is, to have a high portion of signal events in the SR (i.e., $S(\mathcal{X}_s)$) compared to the size of the SR measured by the portion of $4b$ events in the SR (i.e., $P_{4b}(\mathcal{X}_s)$). In an ideal case, the curve should rapidly increase to 1 as the size of SR increases. As Figure 3 suggests, our method can find an SR that shows such behavior, especially when the signal ratio is high (left) or the sample size is large (center). Furthermore, the plot on the right shows that setting $\eta = 0.1$ gives the best performance, demonstrating the importance of choosing a convolution kernel whose size is aligned with the localized signal.

Figure 4 compares our method to a baseline SR proposed by Bryant [3]. The baseline SR is designed to capture events with invariant masses of both intermediate particles close to that of the Higgs boson. We emphasize that the baseline is aware of particular domain knowledge about the signal, which we do not require. While Figure 4 shows a performance gap between the two methods, our model-agnostic SR remains competitive, or even slightly better when the size of SR ($P_{4b}(\mathcal{X}_s)$) is small. In hypothesis testing, it is crucial to limit the size of the SR to maintain a high signal-to-background ratio

¹The rest of the samples were left for the hypothesis testing.




Figure 3: Concentration of signal events in the SR for the $HH \rightarrow 4b$ data. x axis represents the proportion of $4b$ events in the SR (measured by $P_{4b}(\mathcal{X}_s)$) and y axis represents the proportion of signal events in the SR (measured by $S(\mathcal{X}_s)$). As defined previously, n , ϵ , and η are the number of $3b$ events, the signal ratio, and the scale of the convolution kernel, respectively. Bold lines and shaded areas represent (mean) \pm (standard deviation) measured through 10 repeated experiments.




Figure 4: Comparison with the baseline SR provided by Bryant [3]. Bold lines represent the average measured through 10 repeated experiments, and the shaded region represents the standard deviation of our method (the baseline has an ignorable variance, which is hence omitted). While the baseline, which has access to a priori knowledge of the location of the signal, shows better performance when $P_{4b}(\mathcal{X}_s)$ is large, the methods have comparable performance when $P_{4b}(\mathcal{X}_s)$ is small.

in the region. From this perspective, our method achieves signal concentration that is comparable with the baseline while requiring weaker domain-specific knowledge about the signal.

5 Discussion

The plot on the center of Figure 3 shows that an increase in the sample size leads to a better concentration of signals in the SR, even when the signal ratio is fixed to a low value. This suggests that our method has sensitivity for small signals as long as the sample size is large enough. We also remark on the importance of choosing a convolution kernel introducing an optimal amount of noise to the representations of events in learning the smoothed density ratio $\tilde{\gamma}$. If we choose the noise with too large a scale compared to the size of a signal peak, the density ratio would be overly smoothed, affecting the low-frequency features that we intend to preserve in $\tilde{\gamma}$. Conversely, using too small noise would render $\tilde{\gamma}$ indistinguishable from the original density ratio γ . In both cases, we fail to capture the peak we are searching for.

The current scope of our work does not include the estimation of the background $4b$ distribution and the hypothesis testing (Equation (2)) to determine the presence of the signals in the SR. We plan to extend our work to include these steps in future work.

Acknowledgments and Disclosure of Funding

This work was supported in part by NSF grant DMS-2053804. We are grateful to the members of the STAMPS@CMU Research Center for helpful discussions and feedback on this work.

References

- [1] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1):1–29, 2012.
- [2] Blaž Bortolato, Aleks Smolkovič, Barry M. Dillon, and Jernej F. Kamenik. Bump hunting in latent space. *Physical Review D*, 105(11):115009, 2022.
- [3] Patrick Bryant. *Search for Pair Production of Higgs Bosons in the Four Bottom Quark Final State Using Proton-Proton Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector*. PhD thesis, University of Chicago, 2018.
- [4] Purvasha Chakravarti, Mikael Kuusela, Jing Lei, and Larry Wasserman. Model-independent detection of new physics signals using interpretable semi-supervised classifier tests. *The Annals of Applied Statistics*, 17(4):2759–2795, 2023.
- [5] Georgios Choudalakis. On hypothesis testing, trials factor, hypertexts and the BumpHunter. In *PHYSTAT 2011*, 1 2011.
- [6] CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716(1):30–61, 2012.
- [7] CMS Collaboration. Search for Higgs Boson Pair Production in the Four b Quark Final State in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV. *Physical Review Letters*, 129(8):081802, 2022.
- [8] CDF Collaboration. Measurement of $\sigma B(W \rightarrow e\nu)$ and $\sigma B(Z^0 \rightarrow e^+e^-)$ in $\bar{p}p$ collisions at $\sqrt{s} = 1800$ GeV. *Physical Review D*, 44(1):29–52, 1991.
- [9] Jack Collins, Kiel Howe, and Benjamin Nachman. Anomaly Detection for Resonant New Physics with Machine Learning. *Physical Review Letters*, 121(24):241803, 2018.
- [10] Jack H. Collins, Kiel Howe, and Benjamin Nachman. Extending the search for new resonances with machine learning. *Physical Review D*, 99(1):014038, 2019.
- [11] Raffaele Tito D’Agnolo and Andrea Wulzer. Learning new physics from a machine. *Physical Review D*, 99(1):015014, 2019.
- [12] Raffaele Tito D’Agnolo, Gaia Grosso, Maurizio Pierini, Andrea Wulzer, and Marco Zanetti. Learning multivariate new physics. *The European Physical Journal C*, 81(1):1–21, 2021.
- [13] Raffaele Tito D’Agnolo, Gaia Grosso, Maurizio Pierini, Andrea Wulzer, and Marco Zanetti. Learning new physics from an imperfect machine. *The European Physical Journal C*, 82(275): 1–37, 2022.
- [14] Gregor Kasieczka, Benjamin Nachman, David Shih, Oz Amram, Anders Andreassen, Kees Benkendorfer, Blaž Bortolato, Gustaaf Brooijmans, Florencia Canelli, Jack H. Collins, Biwei Dai, Felipe F. De Freitas, Barry M. Dillon, Ioan-Mihail Dinu, Zhongtian Dong, Julien Donini, Javier Duarte, D. A. Faroughy, Julia Gonski, Philip Harris, Alan Kahn, Jernej F. Kamenik, Charanjit K. Khosa, Patrick Komiske, Luc Le Pottier, Pablo Martín-Ramiro, Andrej Matevc, Eric Metodiev, Vinicius Mikuni, Inês Ochoa, Sang Eon Park, Maurizio Pierini, Dylan Rankin, Veronica Sanz, Nilai Sarda, Urous Seljak, Aleks Smolkovic, George Stein, Cristina Mantilla Suarez, Manuel Szewc, Jesse Thaler, Steven Tsan, Silviu-Marian Udrescu, Louis Vaslin, Jean-Roch Vlimant, Daniel Williams, and Mikael Yunus. The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics. *Reports on Progress in Physics*, 84(12):124201, 2021.
- [15] Tudor Manole, Patrick Bryant, John Alison, Mikael Kuusela, and Larry Wasserman. Background modeling for double Higgs boson production: Density ratios and optimal transport. *The Annals of Applied Statistics*, 18(4):2950 – 2978, 2024.
- [16] Benjamin Nachman and David Shih. Anomaly detection with density estimation. *Physical Review D*, 101(7):075042, 2020.