



# A data-directed paradigm for BSM searches: the bump-hunting example

Sergey Volkovich , Federico De Vito Halevy , Shikma Bressler<sup>a</sup>

Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot, Israel

Received: 7 December 2021 / Accepted: 12 March 2022 / Published online: 27 March 2022  
© The Author(s) 2022

**Abstract** We propose a data-directed paradigm (DDP) to search for new physics. Focusing on the data without using simulations, exclusive selections which exhibit significant deviations from known properties of the standard model can be identified efficiently and marked for further study. Different properties can be exploited with the DDP. Here, the paradigm is demonstrated by combining the promising potential of neural networks (NN) with the common bump-hunting approach. Using the NN, the resource-consuming tasks of background and systematic uncertainty estimation are avoided, allowing rapid testing of many final states with only a minor degradation in the sensitivity to bumps relative to standard analysis methods.

## 1 Introduction

Despite its success in describing the elementary particles and their interactions, the Standard Model (SM) is still incomplete [1]. Many models extending beyond the Standard Model (BSM) have been developed over the years predicting the existence of new resonances. Thus, the search for new resonances, either theoretically-predicted or model-agnostic, is a core strategy for discovery in experimental high energy physics (e.g. recently [2–4]).

With almost no exception,<sup>1</sup> all BSM searches have been conducted following the blind analysis paradigm, in which an enormous amount of time and effort is invested before looking at the data, i.e. on background modeling and systematic uncertainty estimation. These resource-intensive tasks have allowed only a limited region in the space spanned by all observables (“observable-space”) to be explored to-date.

<sup>1</sup> In [5–7], discrepancies between the data and Monte Carlo prediction were searched for in a large variety of final states.

<sup>a</sup>e-mail: [shikma.bressler@cern.ch](mailto:shikma.bressler@cern.ch) (corresponding author)

Indeed, searches typically focus on inclusive final states – dilepton, di-photon, di-jet, etc. – ignoring all other observables and avoiding exclusive selections such as di-lepton + jets, di-jets + missing transverse momentum, di-photon within a  $t\bar{t}$  topology, etc. Moreover, within the studied final states, event selection is usually optimized relative to predefined signal models. So far, no significant indication of BSM physics has been found.

Complimentary to the blind analysis paradigm, we propose a data-directed paradigm (DDP) which begins by efficiently identifying regions of interest in the data. Similarly to [5–7], albeit without using Monte Carlo (MC) simulation, the strategy consists of quickly searching the observable-space for exclusive regions exhibiting a significant deviation from some fundamental SM property. Such regions should be considered for data-directed signal hypotheses and further examined using traditional analysis techniques. Like in [4], no MC simulation is used. Thus, the search is not sensitive to MC mismodelling or limited MC statistics. Given the large number of plausible signals which could manifest in an infinite number of exclusive regions, and moreover, the limited time, manpower and resources at hand, searches like the proposed DDP might provide our best chance at discovering BSM physics.

## 2 A data-directed paradigm

A DDP search can be realized with two key ingredients:

1. A theoretically well-established property of the SM with respect to which deviations can be searched for - here we exploit the fact that within the SM, in absence of resonances, almost any invariant mass distribution is smoothly falling. Other properties of the SM, such as flavour symmetry [8] or forward-backward symmetry could also be

exploited once detector effects are taken into account (as implemented for instance in [9]).

2. An efficient algorithm to scan the observable-space in search for deviations – here we train a deep neural network (NN) to map any invariant mass distribution into a distribution of statistical significance for excesses of events (“bumps”). The latter is known as a “z” distribution and is based on the profile likelihood ratio test for positive signals [10]. Different algorithms should be developed when searching for deviations from other SM properties.


The challenge of bump-hunting is an excellent showcase for a search in the DDP<sup>2</sup>; even a simple implementation achieves good accuracy. As long as the underlying background distribution is smoothly falling, a single trained NN, as described in this letter, can quickly perform statistical inference from many selections of observed data. For example, when adapting it to a narrower mass range, it predicted a maximum significance in agreement with the di-muon results presented in [3] in seconds. Crucially, it avoids the time- and effort-consuming tasks of full background and systematic uncertainty estimation currently carried out for every invariant mass distribution under consideration. This way, a potentially unlimited quantity of exclusive distributions can be scanned and large regions in the observable-space can be covered. Nevertheless, event-by-event optimization for bump enhancement, as studied for instance in [11–13], is left to future work.

Bumps identified using the DDP are likely to be caused by statistical fluctuations. These will disappear when tested with more data. Bumps originating from systematic uncertainties due to detector effects (trigger thresholds, kinematic edges, etc.) should appear in MC simulations as well and can be ruled out. Among the bumps found which do not disappear with added data and do not appear in simulation, the most significant ones should be considered as BSM signal hypotheses and devoted a dedicated analysis. Inevitably, some may be due to mismodelled systematic effects.

### 3 A neural network implementation

The NN we employ is trained in a supervised manner, for which we generate a set of artificial training and testing data. These contain inputs which simulate realistic distributions in observed data (in contrast to individual events, as in [11–13]), and can be further tailored for any given search. Inputs are matched to analytically calculated z distributions as targets. When given an invariant mass distribution, the NN predicts a z distribution which shows where and how likely it is that

<sup>2</sup> In [7], bump-hunting was exploited by comparing data to MC prediction in many different final states.



**Fig. 1** Illustration of the sample generation procedure. **a** A smoothly decaying background curve (orange) is generated over 100 bins. Each bin is assigned a Poisson fluctuation. A signal with a significance relative to the fluctuated background (green) is added to it, producing the observed data (blue). **b** The corresponding significance distribution,  $z$ , is calculated analytically. The left and right axes in both panels show the non-scaled and scaled distributions, respectively

the data contains a bump. Once the NN has been trained, we validate that its predictions are consistent and that its loss value converges. Finally, its predictions are evaluated on the test set and we discuss its performance.

We generate the inputs of the NN as 100 bin histograms of observed events,  $N = B + S$ . These are representative of data with high statistics and dynamic range (the bin width reflects a given detector resolution). The generation process is illustrated in Fig. 1. Each input is composed of a smoothly decaying background curve,  $B$ , to which Poisson fluctuations are added, and a localized Gaussian signal,  $S$ , whose significance is defined relative to the fluctuated background. We calculate bin-by-bin the corresponding NN target,  $z$ , which we use to approximate the significance distribution given the unfluctuated background and assuming a Gaussian signal shape [10]. Each input and target pair is collectively referred to as a “sample”. All samples are globally scaled to the interval  $[0, 1]$  under a linear transformation before being utilized by the NN.

A variety of smoothly falling backgrounds is modelled by randomly selecting one of the following ten functional forms for each sample:

$$\begin{aligned}
 & be^{-ax}, \quad ax + b, \quad \frac{1}{ax} + b, \quad \frac{1}{ax^2} + b, \quad \frac{1}{ax^3} + b, \\
 & \frac{1}{ax^4} + b, \quad a(x - x_2)^2 + y_2, \quad -a \cdot \ln(x) + b, \\
 & (y_1 - y_2) \cos(a(x - b)) + y_2, \quad \cosh(a(x - x_2)) + b.
 \end{aligned}
 \tag{1}$$

The parameters  $a$  and  $b$  are defined such that each curve decays between two points,  $(x_1, y_1)$  and  $(x_2, y_2)$ , where  $x_1 < x_2$  are the centers of the extreme bins and  $y_1 > y_2$  are randomized from the interval  $[100, 10,000]$ .

Gaussian shaped signals are generated with mean values distributed randomly between bin 25 and bin 76. The width (standard deviation) of the signals is fixed at 3 bins. To improve the desired feature detection, the NN is trained with a data set containing signals with significance in the range  $[1, 20]\sigma$ . The performance of the NN is determined on a testing data set by evaluating its ability to identify bumps with a significance of  $3\sigma$  – the common definition of a “hint” for BSM physics.


Various NN architectures can be used. Here, we choose an architecture based on a dense layer followed by six 1-dimensional convolutional layers. The latter are intended for feature-detection, while the former is useful in suppressing position-dependent biases. A “rectified linear unit” activation function is used. The “Adam” optimizer is used to minimize the “mean squared error” loss function over 200 epochs at a learning rate of 0.0003 with a batch size of 100. We generate a total of 600,000 training samples, 20% of which are used for validation, and 150,000 testing samples.

### 4 Results

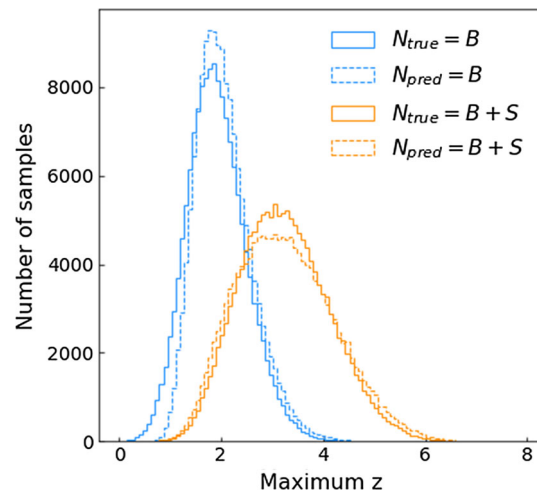
The accuracy of the NN prediction is illustrated in Fig. 2 in terms of the difference between the maximal predicted significance,  $z_{\text{pred}}^{\text{max}}$ , and the one calculated via the profile likelihood ratio test,  $z_{\text{true}}^{\text{max}}$ . All generated test samples are included in the figure; in over 87% of these the predicted peak was found within 1 bin of  $z_{\text{true}}^{\text{max}}$ . A mean ( $\mu$ ) of  $-0.02$  indicates a negligible bias in the prediction and a 0.46 standard deviation ( $\sigma$ ) measures its precision. The asymmetry seen as a sharp edge in the third quadrant originates from the small number of maximal  $z$  predictions below one.

We are interested in finding samples with bumps of  $3\sigma$  significance while rejecting samples without bumps. Figure 3 shows  $z_{\text{true}}^{\text{max}}$  in a solid line and  $z_{\text{pred}}^{\text{max}}$  in a dashed line for samples with no signal added (blue) and for samples with a  $3\sigma$  significance signal added (orange). In a traditional bump-hunting search, the signal significance is evaluated relative to an estimated background. Thus, the measured significance of a  $3\sigma$  signal could fluctuate around this value. This is the origin of the width of the  $z_{\text{true}}^{\text{max}}$  distributions: the signal is generated with a significance relative to the fluctuated background and its  $z_{\text{true}}^{\text{max}}$  is evaluated relative to the smooth background.

According to the Neyman–Pearson lemma (see e.g. [14]),  $z_{\text{true}}^{\text{max}}$  provides the most powerful signal to background separation. It relies on exact knowledge of both the background and signal shapes. Yet, despite using no a priori knowledge of the two, the signal to background separation in  $z_{\text{pred}}^{\text{max}}$  is only




**Fig. 2** The difference between  $z_{\text{pred}}^{\text{max}}$  and  $z_{\text{true}}^{\text{max}}$  as a function of  $z_{\text{true}}^{\text{max}}$ . Dense regions are shown in red (roughly corresponding to the  $1\sigma$  region), while sparse regions are shown in blue



**Fig. 3** The distribution of  $z_{\text{true}}^{\text{max}}$  (solid line) and  $z_{\text{pred}}^{\text{max}}$  (dashed line) for samples with no signal added (blue) and for samples with a  $3\sigma$  significance signal added (orange)

slightly degraded relative to  $z_{\text{true}}^{\text{max}}$ . This is quantified in terms of receiver operating characteristic (ROC) curves in Fig. 4, obtained from the distributions of Fig. 3. The true (blue) and predicted (orange) ROC curves show the efficiency to correctly identify a  $3\sigma$  bump versus the false positive rate of selecting samples with no injected bump. The area under the true curve,  $A_{\text{true}}$ , is 0.899 while the area under the predicted curve,  $A_{\text{pred}}$ , is 0.865, which implies a degradation in performance of less than 4%. In other words, the probability that based on the NN output a selection will be marked as potentially interesting approaches the probability that a traditional method would do the same.

We also confirmed that the NN is able to generalize in identifying with comparable accuracy bumps over linear combi-



**Fig. 4** True (blue) and predicted (orange) ROC curves and their associated areas under curve,  $A_{true}$  and  $A_{pred}$

nations of the background forms (Eq. 1), and over an unseen 10th shape when trained on 9 background shapes.<sup>3</sup> Thus, it is unrestricted by specific background forms in its capacity to detect bumps, which goes beyond the potential of traditional techniques. Similar performance was obtained in additional scenarios: when testing on distributions with lower and higher statistics (in the ranges between 100–500 and 5000–10,000, respectively), when extending the bump region from the bin range 25–76 to 5–96, and when training and testing on samples with wider bumps, of either 4 or 5 bins.

## 5 Validation

We validate the convergence of the loss value achieved by increasing either the number of epochs or the size of the training data set. In terms of  $A_{pred}$  from Fig. 4, the NN performance varies insignificantly, by less than 1% when moving past 200 epochs (for 100,000 input samples) or 500,000 input samples (for 100 epochs).

Consistency was ensured by comparing the NN predictions in two scenarios (with 100,000 training samples and 100 epochs). First, we trained four different NNs using an independent training data set for each and compared their predictions on a common testing data set (with 25,000 samples). Second, the performance of each of the NNs was separately compared on four different testing data sets. In all cases, the accuracy when separating signal from background was unaffected.

<sup>3</sup> This test was carried out for the background shapes  $be^{-ax}$  and  $\frac{1}{ax^4} + b$ .

## 6 Discussion

We have presented a data-directed paradigm, complementary to the blind analysis paradigm, and demonstrated one of its possible implementations using the concept of bump-hunting. We have shown that a NN can be trained to efficiently identify bumps over smoothly falling backgrounds without being given any a priori information about the background or the bump's position. Relative to the most powerful test statistic (profile likelihood ratio), which relies on exact knowledge of both the background and signal shapes, the performance of the NN was only inferior by less than 4% when considering the area under the ROC curve. Since for each data distribution the NN prediction is obtained within a couple of seconds (compared to a year or more when following the blind analysis paradigm), these results pave the way towards scanning the overwhelming observable-space that is being measured in experiments searching for bumps. Examples could be searches for di-lepton, di-jet, di-photon, or jet-lepton-missing  $E_T$  resonances, in events containing, in addition, any other set of objects.

In the search for BSM physics we must leave no stone unturned. Complementary to traditional theory-directed blind analysis searches, the DDP should be pursued as well. With the expected ramp up of the Large Hadron Collider, existing data should be thoroughly explored. A first milestone could be demonstrating sensitivity to bumps in regions already investigated. If needed, dedicated NNs could be trained to account for scenarios not covered by the current implementation (e.g. different dynamic ranges, bins or widths) and other architectures could be explored. The search for bumps over smoothly falling backgrounds is just one example of a property of the SM that could be considered. Others such as flavour symmetry [8] or forward-backward symmetry could be exploited as well. Given the challenge ahead, searches like the proposed DDP might provide our best chance at discovering BSM physics.

**Funding** This work was supported by Grant no. 2871/19 from the Israeli Science Foundation (ISF), Grant no. I-1506-303.7/2019 from the German Israeli Foundation (GIF) and by the Sir Charles Clore Prize.

### Declarations

**Conflict of interest** The authors have no conflicts of interest to declare that are relevant to the content of this article.

**Code availability** The code of this work is available from the corresponding author on reasonable request.

**Availability of data and material** The data sets generated and analysed during this work are available from the corresponding author on reasonable request.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.  
Funded by SCOAP<sup>3</sup>.

## References

1. S. Weinberg, Phys. Rev. Lett. **121**, 220001 (2018)
2. ATLAS Collaboration, JHEP **03**, 145 (2020)
3. ATLAS Collaboration, Phys. Lett. B **796**, 68-87 (2019)
4. ATLAS Collaboration, Phys. Rev. Lett. **125**, 051801 (2020)
5. ATLAS Collaboration, Eur. Phys. J. C **79**, 120 (2019)
6. CMS Collaboration, Eur. Phys. J. C **81**, 629 (2021)
7. CDF Collaboration, Phys. Rev. D **79**, 011101 (2009)
8. S. Bressler, A. Dery, A. Efrati, Phys. Rev. D **90**, 015025 (2014)
9. ATLAS Collaboration, Eur. Phys. J. C **77**, 70 (2017)
10. G. Cowan, K. Cranmer, E. Gross, O. Vitells, Eur. Phys. J. C **71**, 1554 (2011)
11. M. Farina, Y. Nakai, D. Shih, Phys. Rev. D **101**, 075021 (2020)
12. J. Collins, K. Howe, B. Nachman, Phys. Rev. D **99**, 014038 (2019)
13. J. Collins, K. Howe, B. Nachman, Phys. Rev. Lett. **121**, 241803 (2018)
14. E.L. Lehmann, J.P. Romano, *Testing Statistical Hypotheses*, 3rd edn. (Springer, New York, 2005)