Beyond $M_{t\bar{t}}$: learning to search for a broad $t\bar{t}$ resonance at the LHC

Sunghoon Jung^a, Dongsub Lee^b, Ke-Pan Xie^c

Center for Theoretical Physics, Department of Physics and Astronomy, Seoul National University, Seoul 08826, South Korea

Received: 2 August 2019 / Accepted: 25 January 2020 / Published online: 8 February 2020
© The Author(s) 2020

Abstract A resonance peak in the invariant mass spectrum has been the main feature of a particle at collider experiments. However, broad resonances not exhibiting such a sharp peak are generically predicted in new physics models beyond the Standard Model. Without a peak, how do we discover a broad resonance at colliders? We use machine learning technique to explore answers beyond common knowledge. We learn that, by applying deep neural network to the case of a $t\bar{t}$ resonance, the invariant mass $M_{t\bar{t}}$ is still useful, but additional information from off-resonance region, angular correlations, p_T , and top jet mass are also significantly important. As a result, the improved LHC sensitivities do not depend strongly on the width. The results may also imply that the additional information can be used to improve narrow-resonance searches too. Further, we also detail how we assess machine-learned information.

1 Introduction

Discovering new physics through a new resonance is one of the most exciting opportunities. A “narrow” resonance peak, being sharply localized in the energy spectrum, allows for the most efficient discovery above continuum backgrounds as well as for precision measurements of the particle mass, width and other properties. However, the widths of new (and presumably heavier) resonances in new physics can be easily much larger than those of the Standard Model (SM) particles. The width generally grows with the mass of a resonance, and a new strong coupling may induce rapid decay as in composite Higgs models [1–4] or warped extra dimensional models [5, 6]. Also, more decay channels to lighter beyond-SM particles may open up, which further increases the width.

The large width causes several difficulties in collider experiments. Above all, without a sharp peak, the discovery becomes challenging, as the signal becomes spread over a large range of energy above continuum backgrounds. For example, the ATLAS result based mostly on the invariant mass distribution [7] shows that for a $M = 1$ TeV Kaluza–Klein gluon, the measured (expected) cross section upper limit $\sigma(pp \rightarrow g_{KK} \rightarrow t\bar{t})$ increases from 1.4 (1.2) pb to 4.7 (2.7) pb when the width-to-mass ratio Γ/M varies from 10 to 40%. In addition, the phenomenological study in Ref. [4] shows that for the minimal composite Higgs model with the third generation left-handed quark $q_L = (t_L, b_L)^T$ being fully composite, a vector $t\bar{t}$ resonance as light as $M = 1$ TeV is still allowed by the direct search in the $\Gamma/M \gtrsim 20\%$ region.

Secondly, broad resonance shape is more susceptible to the energy dependences of parton luminosity and the width, interferences with backgrounds or other resonances, and mixing and overlap with nearby resonances. These effects make discoveries further challenging and complicated. In particular, the complex interference (the one with imaginary parts in amplitudes) in supersymmetric or two-Higgs doublet models can make broad heavy Higgs bosons decaying to $t\bar{t}$ generally appear not as a pure resonance peak [8–11] but even as pure dips or nothing [11]. And nearly degenerate heavy Higgs bosons can overlap significantly, producing complicated resonance shapes [12–14].

Many of these new broad resonances are just beyond the current reach of the LHC. Thus, it is imperative to study the physics of broad resonances and develop efficient discovery methods. However, broad-resonance searches have been studied only in limited cases, e.g., phenomenologically in Refs. [4] (third-generation quark pair, $\ell^+\ell^-$), [5] ($\mu^+\mu^-$), and experimentally in Refs. [7, 15] ($t\bar{t}$), [16, 17] (jj) [18] ($\ell^+\ell^-$). In all the cases, the invariant mass had still been used as a main observable, but the question of “how do

^a e-mail: sunghoonj@snu.ac.kr

^b e-mail: dongsub93@snu.ac.kr

^c e-mail: kpxie@snu.ac.kr

we (best) discover a broad resonance without a peak?” had not been answered thoroughly.¹

This question might be a problem appropriate to use deep neural network (DNN) technique to answer. It is because the answer is not so obvious, a priori, and even small improvements will be significant. Machine learning has indeed been applied to various problems in particle physics. For example, bump-hunting resonance searches were improved with DNN [20, 21]. The DNN is one of machine learning algorithms. Coming with various network structures such as fully-connected network [22–27], convolutional neural network [28–31] and others [32–36], DNN had shown remarkable performances in the exploration of physics beyond the SM, often better than other machine learning algorithms such as boosted decision tree (BDT). We refer to Refs. [37, 38] and references therein for reviews of the DNN applications in LHC physics.

In this paper, we consider a spin-1 broad $t\bar{t}$ resonance at the LHC (Sect. 2). Being the heaviest particle in the SM, the top quark has been regarded as an important portal to new physics. As a first step toward a more general study of broad resonances, we ignore any interference effects and nearby resonances (Sect. 3). We use fully-connected DNN to explore answers beyond common knowledge (Sect. 3). Finally, we assess whether and what DNN can learn, even beyond what we know well (Sect. 4).

2 Benchmark model

For simplicity, here we consider a gauge singlet vector resonance ρ interacting strongly with the SM right-handed top quark t_R , and the relevant Lagrangian is

$$\mathcal{L} = -\frac{1}{4}\rho_{\mu\nu}\rho^{\mu\nu} + \frac{m_\rho^2}{2g_\rho^2}(g_\rho\rho_\mu - g_1 B_\mu)^2 + \bar{t}_R \gamma^\mu t_R (g_\rho\rho_\mu - g_1 B_\mu), \quad (1)$$

where $\rho_{\mu\nu} = \partial_\mu\rho_\nu - \partial_\nu\rho_\mu$, and g_1 is the SM hypercharge gauge coupling. This model is also considered in Refs. [2, 19, 39]. Note that the ρ_μ mixes with the SM gauge field B_μ in Eq. (1). Given $g_\rho \gg g_1$, the mixing angle is $\sin\theta \approx g_1/g_\rho$ before the electroweak symmetry breaking (EWSB). Therefore, after transforming to the mass eigenstates, the interactions between ρ resonance and SM fermions will be $\sim g_\rho$ for t_R and $\sim Y g_1^2/g_\rho$ for other fermions (including t_L and other light quarks), with Y being the hypercharge of the corresponding fermion. The physical mass of ρ is $M_\rho = m_\rho$. EWSB gives $\mathcal{O}(v^2/m_\rho^2)$ corrections to above picture, and the details can be found in Appendix C of Ref. [19].

¹ If a broad resonance can decay to multi-top/ W channels, it can be searched using the same-sign di-lepton final state [3, 4, 19], which doesn't rely on the reconstruction of the invariant mass.

Due to the large coupling g_ρ , the ρ resonance decays to $t\bar{t}$ with a branching ratio $\sim 100\%$, and the width-to-mass ratio is

$$\frac{\Gamma_\rho}{M_\rho} \approx \frac{\Gamma_{\rho \rightarrow t\bar{t}}}{M_\rho} \approx \frac{g_\rho^2}{8\pi}. \quad (2)$$

For $g_\rho = 3$ and 4, this ratio reaches 36% and 64%, respectively. Thus a broad ρ is easily realized in the model described by Eq. (1). Note that $\Gamma_{\rho \rightarrow t\bar{t}} \leq \Gamma_\rho$, if ρ has other strong dynamical decay channels such as the decay to low-mass top partners (which are not listed in our simplified model), typically Γ_ρ is several times larger than $\Gamma_{\rho \rightarrow t\bar{t}}$, thus a large Γ_ρ/M_ρ can be obtained even for smaller g_ρ . We consider $M_\rho = 1$ and 5 TeV as two benchmarks, and for each mass point the width-to-mass ratios $\Gamma_\rho/M_\rho = 10\%, 20\%, 30\%$ and 40% are considered. The corresponding benchmark cases are then identified as $Mi\Gamma j$, with $i = 1$ or 5 denoting the mass (in unit of TeV) and $j = 1, 2, 3, 4$ being $10 \times \Gamma_\rho/M_\rho$. For example, M1 Γ 4 is the benchmark for $M_\rho = 1$ TeV and $\Gamma_\rho/M_\rho = 40\%$.

At the LHC, the ρ resonance can be produced via the Drell–Yan process ($q\bar{q} \rightarrow \rho$) through the ρ -light quark interaction. Among the various decay channels of the $t\bar{t}$, we choose to focus on the semi-leptonic final state

$$pp \rightarrow \rho \rightarrow t\bar{t} \rightarrow \ell^\pm \nu b\bar{b}jj. \quad (3)$$

The dominant background is then the SM $t\bar{t}$ process, which contributes 81% $\sim$ 88% of the total backgrounds [7]. For simplicity, we only consider this background. It should be emphasized that although we provide a benchmark model as physical motivation here, our results are general for all heavy singlet spin-1 resonances with top quark portal.

3 Searching for a broad $t\bar{t}$ resonance

In this section, we describe technical details of our work and show final cross section limits. First, we describe how we parameterize a broad resonance, and how we build learning datasets and train DNN for each benchmark signal case. Then we derive improved cross section upper limits.

3.1 Breit–Wigner description

We assume a single, isolated broad resonance far away from any other resonances and thresholds, and ignore any interference effects. Then we use the following Breit–Wigner description of the propagator of a broad resonance

$$\frac{1}{s - \hat{M}_\rho^2(s) + i\sqrt{s}\hat{\Gamma}_\rho(s)} \approx \frac{1}{s - M_\rho^2 + iM_\rho\Gamma_\rho}, \quad (4)$$

where the nominal resonance mass M_ρ and the width Γ_ρ are fixed constants. The energy dependence of the mass $\hat{M}_\rho(s)$ from the real part of the self-energy correction is higher-order, hence small irrespective of the large width. On the other hand, the energy dependence of the width $\hat{\Gamma}_\rho(s) \propto \sqrt{s}$ from the imaginary part can induce corrections as large as ~ 100 (10)% for broad resonances considered in this paper $\Gamma_\rho/M_\rho \sim 40$ (20)%. But, within this range of the width, the resonance shape remains relatively undistorted albeit some shifts of the peak and height [4, 5, 40]. Also, the fixed mass and width have been used in LHC searches of broad resonances [7, 15]. Thus, we use Eq. (4) with fixed M_ρ and Γ_ρ , both for simplicity and for comparison purpose.

3.2 Preparing training data

The model described by Eq. (1) is written in the universal FeynRules output file [41]. We generate parton-level events of the signals and background using 5-flavor scheme within the MadGraph5_aMC@NLO [42] package. All spin correlations of the final state $\ell^\pm \nu b \bar{b} j j$ objects are kept. The phase space integrate region is set to $|\sqrt{s} - M_\rho| \leq 15 \times \Gamma_\rho$, which is large enough for us to simulate the full on- and off-shell effect of the ρ resonance. The interference between $pp \rightarrow \rho \rightarrow t\bar{t}$ and the SM $t\bar{t}$ background is negligible [7], thus not considered here. We normalize the SM $t\bar{t}$ cross section with the next-to-next-to-leading order with next-to-next-to-leading logarithmic soft-gluon resummation calculation from the Top++2.0 package [43–48], and the K -factor is 1.63. The parton-level events are matched to +1 jet final state and then interfaced to Pythia 8 [49] and Delphes [50] for parton shower and fast detector simulation. As for the detector setup, we mainly use the CMS configuration, but with following modifications: the isolation ΔR parameters for electron, muon and jet are set to 0.2, 0.3 and 0.5 respectively. The b -tagging efficiency (and mis-tag rate for c -jet, light-flavor jets) is corrected to 0.77 (and 1/6, 1/134) according to Ref. [51]. We generate $\gtrsim 5 \times 10^6$ events for the background and each signal benchmark.

We defined two kinematic regions. The first one is called the resolved region, in which the decay products of the top quark (i.e. $\ell^\pm \nu b \bar{b} j j$) are identified as individual objects. This region is defined as follows

1. Exactly one charged lepton $\ell^\pm = e^\pm$ or $\mu^\pm$ with $p_T^\ell > 30$ GeV and $|\eta^\ell| < 2.5$. Events containing a second lepton with $p_T^\ell > 25$ GeV are vetoed.
2. $\cancel{E}_T > 20$ GeV and $\cancel{E}_T + M_T^W > 60$ GeV, where the W -transverse mass is defined as

$$M_T^W = \sqrt{2p_T^\ell \cancel{E}_T [1 - \cos \Delta\phi(p_T^\ell, \cancel{E}_T)]}.$$

3. At least four jets with $p_T^j > 25$ GeV and $|\eta^j| < 2.5$, and at least one of the leading four jets is b -tagged.

The cuts are mainly based on Ref. [7], but with some simplifications. The cut flows of the signals and backgrounds are listed in Table 1. We only consider the $M_\rho = 1$ TeV benchmark cases in this kinematic region. The SM $t\bar{t}$ cross section is 68.9 pb taken into account the K -factor.

The second kinematic region is the boosted region, in which the hadronic decay products of the top quark are combined into a fat jet. The corresponding event selection criteria is

1. Exactly one charged lepton $\ell^\pm = e^\pm$ or $\mu^\pm$ with $p_T^\ell > 30$ GeV and $|\eta^\ell| < 2.5$. Events containing a second lepton with $p_T^\ell > 25$ GeV are vetoed.
2. $\cancel{E}_T > 20$ GeV and $\cancel{E}_T + M_T^W > 60$ GeV.
3. Exactly one top-jet with $p_T^{j_{\text{top}}} > 300$ GeV and $|\eta^{j_{\text{top}}}| < 2.0$, and satisfies $\Delta\phi(j_{\text{top}}, \ell^\pm) > 2.3$. The top-jet is reconstructed with a $R = 1.0$ cone in anti- k_t algorithm, and is trimmed with $R_{\text{cut}} = 0.2$ and $f_{\text{cut}} = 0.05$ [52]. We use a simplified top-tagging procedure in event selection. The top-tagging efficiency and the mistag-rate are set to 80% and 20% respectively, based on Ref. [53], which makes use of jet invariant mass and N -subjettiness [54–59].
4. Exactly one selected jet with $p_T^{j_{\text{sel}}} > 25$ GeV and $|\eta^{j_{\text{sel}}}| < 2.5$. In addition, the selected jet should have $\Delta R(j_{\text{sel}}, j_{\text{top}}) > 1.5$ and $\Delta R(j_{\text{sel}}, \ell) < 1.5$.

The cuts here are again mainly based on Ref. [7], and the cut flows for signals and background are listed in Table 2. In this region, we consider both $M_\rho = 1$ and 5 TeV signals. To increase the event generating efficiency of the background events, in this region we require the SM $pp \rightarrow t\bar{t} \rightarrow \ell^\pm \nu b \bar{b} j j$ process has at least one final state parton (including the b -parton) with $p_T > 150$ GeV. This is done by setting `xptj = 150` in MadGraph5_aMC@NLO. We have checked that this setup doesn't lose the generality, but improves the event generating efficiency by a factor of ~ 6 . The background cross section after cuts is 2.88 pb taken into account the K -factor.

The events after cuts are collected to make training and validation/test datasets. For the resolved region, we have $1\ell^\pm + \cancel{E}_T + 4$ jets in total 6 reconstructed objects in the final state, and 26 low-level kinematic observables can be used as input features: E^ℓ , p_T^ℓ , η^ℓ and ϕ^ℓ from the charged lepton; $\cancel{E}_T$, $\phi^{\cancel{E}_T}$ from the missing transverse momentum; E^{j_i} , $p_T^{j_i}$, η^{j_i} , ϕ^{j_i} and b^{j_i} from the 4 leading jets, with $i = 1, 2, 3, 4$. Here b^j is the b -tagging observable, which is 1 for a b -tagged jet and 0 otherwise. Some examples of the low-level observables distributions are shown in Fig. 1a. For

Table 1 The cut flows of the signals and background in resolved region. The events are generated at $1\ell^\pm + \cancel{E}_T + \text{jets}$ final state, where ℓ denotes e and μ . $M_i\Gamma j$ denotes the benchmark case with $M_\rho = i$ TeV and $\Gamma_\rho/M_\rho = 0.1 \times j$

Process	Event number	Cut 1	Cut 2	Cut 3	Efficiency
$M1\Gamma1$	5.00×10^6	3.32×10^6	3.02×10^6	1.81×10^6	36.3%
$M1\Gamma2$	5.00×10^6	3.29×10^6	2.98×10^6	1.79×10^6	35.8%
$M1\Gamma3$	3.85×10^6	2.52×10^6	2.23×10^6	1.36×10^6	35.3%
$M1\Gamma4$	5.00×10^6	3.25×10^6	2.93×10^6	1.75×10^6	34.9%
$SM\ t\bar{t}$	4.98×10^6	2.60×10^6	2.21×10^6	1.39×10^6	28.0%

Table 2 The cut flows of the signals and background in boosted region. The events are generated at $1\ell^\pm + \cancel{E}_T + \text{jets}$ final state, where ℓ denotes e and μ . $M_i\Gamma j$ denotes the benchmark case with $M_\rho = i$ TeV and

$\Gamma_\rho/M_\rho = 0.1 \times j$. The setup $x_{\text{ptj}} = 150$ is to improve the event generating efficiency of the background, see the text for details

Process	Event number	Cut 1	Cut 2	Cut 3	Cut 4	Efficiency (%)
$M1\Gamma1$	5.00×10^6	3.32×10^6	3.02×10^6	1.17×10^6	9.61×10^5	19.2
$M1\Gamma2$	5.00×10^6	3.29×10^6	2.98×10^6	1.08×10^6	8.85×10^5	17.7
$M1\Gamma3$	5.00×10^6	3.27×10^6	2.96×10^6	1.02×10^6	8.34×10^5	16.7
$M1\Gamma4$	5.05×10^6	3.28×10^6	2.96×10^6	9.92×10^5	8.06×10^5	15.9
$M5\Gamma1$	5.00×10^6	2.53×10^6	2.36×10^6	1.15×10^6	8.41×10^5	16.8
$M5\Gamma2$	5.00×10^6	2.72×10^6	2.52×10^6	1.19×10^6	8.76×10^5	17.5
$M5\Gamma3$	5.00×10^6	2.81×10^6	2.59×10^6	1.19×10^6	8.85×10^5	17.7
$M5\Gamma4$	5.00×10^6	2.86×10^6	2.64×10^6	1.20×10^6	8.90×10^5	17.8
$SM\ t\bar{t} (x_{\text{ptj}} = 150)$	1.99×10^7	1.22×10^7	1.08×10^7	1.41×10^6	1.21×10^6	6.10

each benchmark case (i.e. $M1\Gamma1 \sim M1\Gamma4$), we build a training dataset and a validation/test dataset. Both of those two datasets have 1,000,000 events, which contain nearly equal signal and background events.

For the boosted region, $1\ell^\pm + \cancel{E}_T + 1 \text{ top-jet} + 1 \text{ selected jet}$ in total 4 objects are reconstructed, and we can extract 15 low-level observables as input features: the first 6 are from ℓ and $\cancel{E}_T$, same as the resolved region; the other 9 consist of $E_{\text{T}}^{j_{\text{sel}}}$, $p_{\text{T}}^{j_{\text{sel}}}$, $\eta^{j_{\text{sel}}}$, $\phi^{j_{\text{sel}}}$, $b^{j_{\text{sel}}}$ from the selected jet, and $E_{\text{T}}^{j_{\text{top}}}$, $p_{\text{T}}^{j_{\text{top}}}$, $\eta^{j_{\text{top}}}$, $\phi^{j_{\text{top}}}$ from the top-jet. Some examples of the low-level observables distributions are illustrated in Fig. 1b. For each benchmark case (i.e. $M1\Gamma1 \sim M1\Gamma4$, and $M5\Gamma1 \sim M5\Gamma4$), we randomly mix equal number of signal and background events to get 800,000 events for training and another 800,000 events for validation/test.

3.3 Training the DNN

The DNN classifier is implemented using the Keras [60] package (with Tensorflow [61] as the backend). The architecture of the DNN is as follows,

$$[N_{\text{in}}, \underbrace{N_{\text{node}}, N_{\text{node}}, \dots, N_{\text{node}}}_{N_{\text{hidden}}}, 2], \quad (5)$$

where N_{hidden} and N_{node} are the numbers of hidden layers and the number of neurons per hidden layer, respectively. The number of input features $N_{\text{in}} = 26$ (15) for the resolved

(boosted) region. All the input features are rescaled to have average 0 and standard deviation 1 before training. We label the events with column matrices to match the two neurons in output layer:

$$\text{signal} \rightarrow \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad \text{background} \rightarrow \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \quad (6)$$

The Rectified Linear Unit (ReLU) activation function is used for all the hidden layers, while the softmax activation function is adopted for the output layer. The loss function is categorical_crossentropy, and the optimizer is Adam. To get the best configuration of the DNN, we try various choices of the hyper-parameter combination as follows,

$$\begin{aligned} N_{\text{hidden}} &= 4, 5; & N_{\text{node}} &= 200, 300; \\ L_r &= 0.001, 0.003; & D_r &= 0.1, 0.2, 0.3; \\ B_s &= 10^3, 10^4; \end{aligned} \quad (7)$$

where L_r is the initial learning rate, D_r is the dropout rate, and B_s is the batch size. For each benchmark case, there are in total 48 different DNN configurations, in which we select the best one based on the learning curves with the following criteria:

1. If the validation/test accuracy curve achieves its maximum when crossing with the training accuracy curve, and meanwhile the validation/test loss curve reaches its

Fig. 1 Distributions of some low-level observables used to train a DNN. Selection cuts in Tables 1 and 2 are applied

minimum and crosses with the training curve, we select that configuration and cut the training at that epoch. This early stop is to prevent over-fitting.

2. If more than one configurations have the behaviors mentioned above, then we select the one with the higher validation/test accuracy and lower validation/test loss; if still there remain more than one networks, we choose the one with learning curves having less fluctuation.

Fig. 2 The final DNN output r distributions that we use to obtain cross section upper limits. Benchmark cases with $\Gamma_\rho/M_\rho = 40\%$ and $t\bar{t}$ backgrounds

The details of training and the chosen configurations are listed in Tables 3 and 4 of the Appendix. For the $M_\rho = 1$ TeV models, the DNN can reach a classification accuracy of $\geq 80\%$ in the resolved region and of $\geq 65\%$ in the boosted region. While for the $M_\rho = 5$ TeV case, the accuracy is $\geq 76\%$ in the boosted region.

The softmax activation function for the output layer guarantees the output responses of the 0th neuron (r_0) and the 1st neuron (r_1) satisfy

$$0 < r_0, r_1 < 1, \quad r_0 + r_1 \equiv 1. \quad (8)$$

Therefore, we can consider r_1 only, and denote it as r . Due to the label definition in Eq. (6), If the DNN is well trained, the distribution of r should have a peak around 1.0 (0.0) for the signal (background), for both the training data and the validation/test data. Figure 2 shows the distributions of the validation/test data for benchmark cases with $\Gamma_\rho/M_\rho = 40\%$ as an illustration. The DNN for M1 Γ 4 shows worse performance in boosted region compare to the one in resolved region. This is because that two peaks in neuron output from signal and SM background are not separated well. In fact, this is a generic feature for all $M_\rho = 1$ TeV benchmark cases. It is mainly due to the boosted region cuts, which require a top-jet with $p_T^{J1op} > 300$ GeV. As a result, most of the SM $t\bar{t}$ background events are round this value. However, for a $M_\rho = 1$ TeV resonance, its decay product $t/\bar{t}$ acquires a transverse momentum ~ 500 GeV, quite similar to the cut threshold. Therefore, the signal and background look similar

(see the $p_T^{j_{\text{top}}}$ distribution in Fig. 1b), and thus the separation is not efficient. On the other hand, for a $M_\rho = 5$ TeV resonance, $p_T^{j_{\text{top}}} \sim 2.5$ TeV, the DNN works very well, as plotted in the bottom of Fig. 2.

3.4 Setting bounds for the signal

We treat the neuron output r as an observable, and fit its distribution shape to get the cross section upper limit of $pp \rightarrow \rho \rightarrow t\bar{t}$ for a given integrated luminosity. For the $M_\rho = 1$ TeV benchmark cases, we use a binned χ^2 fitting method by dividing the $0 < r < 1$ range into 50 bins. While for the $M_\rho = 5$ TeV benchmarks, as the signal cross sections are expected to be tiny, to improve the efficiency we use the un-binned fitting method described in Refs. [62, 63]. In each case, we consider the statistic uncertainty and assume a 12% systematic uncertainty for the background. To include the effect of other subdominant backgrounds besides $t\bar{t}$ (i.e. $W + \text{jets}$, multi-jet, etc), we further rescale the cross section by a factor of $1.23 = 1/0.81$ and $1.14 = 1/0.88$ for the resolved and boosted regions, respectively. Those factors come from the fact that $t\bar{t}$ contributes 81% (88%) of the total background for resolved (boosted) region [7]. This simple rescaling could overestimate final contributions from subdominant backgrounds, and result in somewhat conservative estimations of cross section bounds.

The signal strength upper limits are derived for the unfolded parton-level cross section $\sigma(pp \rightarrow \rho \rightarrow t\bar{t})$, which can be compared with the final results in experimental papers, e.g. Refs. [7, 15]. Our results are shown in Fig 3, in which the expected and measured upper limits of Ref. [7] are also plotted as references, as they use the same final state and similar selection cuts. One can read that the DNN results are rather insensitive to the width of the ρ resonance compare to the traditional approach, achieving better constraints in the large width region.² For the $M_\rho = 1$ TeV benchmark, the result is obtained by the combined fitting of both resolved and boosted regions. Individually, the resolved and boosted regions respectively yield cross sections ~ 3 pb and ~ 1 pb. Although networks in the resolved region have a higher accuracy ($\geq 80\%$) than those in the boosted region ($\geq 65\%$) in Table 3, they actually give a worse measurement of the cross section. This is because the boosted cuts can remove lots of background events and hence improve the fitting performance. That is also the reason why we only consider the boosted region for $M_\rho = 5$ TeV: the production rate for such a high mass ρ is so small that we have to use the boosted region to suppress the background. The DNN bounds for 5 TeV signal benchmark are comparable to the experimentally measured ones, but still better than the experimentally

expected ones. As the training uses random number for the initialization of weights and biases, even for a given DNN configuration, the final results are slightly different for different running. To take into account this training uncertainty, we repeat 15 times of running the chosen DNN configuration for each benchmark case. For the $M_\rho = 1$ TeV case, the relative fluctuation is small thus not shown; while for the $M_\rho = 5$ TeV case, the standard deviations of the runs are shown as vertical error bars in Fig. 3.

4 Figuring out what the machine had learned

In this section, we attempt to assess information learned by DNN using three methods, each of which will be discussed in each subsection. As a result, we can figure out not only which information has been learned, but also which information is most important.

4.1 Testing high-level observables

It is important to know whether a DNN had learned well-known useful but complicated features. In fact, it has been argued that some machine learning methods such as jet image [30] do not efficiently capture invariant mass features [29].

Our approach is to train another set of DNNs using additional high-level observables, of which features we want to test. By comparing the performances of these new DNNs with the original DNNs trained with only low-level observables, we can test whether those particular high-level features (i.e. physically-motivated) have been effectively learned³ or not. This “saturation approach” has been widely used in particle physics research [23, 64].

To construct high-level observables, we first reconstruct the t and $\bar{t}$. The longitudinal momentum of the neutrino is solved by requiring the leptonically decaying W to be on-shell, i.e. $M_{\ell\nu} = M_W$. For the resolved region, the assignment of the 4 reconstructed jets are done by minimizing

$$\chi^2 = \frac{(M_{jj} - M_W)^2}{\sigma_W^2} + \frac{(M_{jjj} - M_t)^2}{\sigma_t^2} + \frac{(M_{j\ell\nu} - M_t)^2}{\sigma_t^2},$$

for various jet permutations, where $\sigma_W = 0.1 \times M_W$ and $\sigma_t = 0.1 \times M_t$. For the boosted region, a top quark is identified as the top-jet and the other is reconstructed from the combination of $\ell^\pm \nu_{j_{\text{sel}}}$. Once the t and $\bar{t}$ are reconstructed, we are able to define the following 7 high-level observables for the signal $pp \rightarrow \rho \rightarrow t\bar{t}$:

² We also checked that the DNN results are better than those from more traditionally used BDT.

³ The definition of “learned” could be ambiguous, but we use subjective criteria discussed in text.

Fig. 3 The DNN-improved cross section upper limits at 95% C.L., obtained by fitting DNN output r distributions in Fig. 2. The latest ATLAS results [7] are also shown for comparison. The vertical error bars of the DNN results are training uncertainties, which are derived by running the same network for 15 times

1. The invariant mass $M_{t\bar{t}}$ of the $t\bar{t}$ system.
2. The polar angle and azimuthal angle in the Collins-Soper frame [65]. We label the leptonic and hadronic decaying tops with subscripts “tl” and “th”, respectively. Hence we have $\cos\theta_{tl}^{CS}$, $\cos\theta_{th}^{CS}$, ϕ_{tl}^{CS} and ϕ_{th}^{CS} in total 4 observables.
3. The polar angles in the Mustraal frame [66], $\cos\theta_1^{Mus.}$ and $\cos\theta_2^{Mus.}$.

The first observable reveals the resonance feature, while the latter 6 observables reflect the spin-1 nature of the ρ resonance. For the boosted region, to take into account the features of the top jet, we introduce 3 additional high-level observables, i.e.

1. The invariant mass $M_{j_{top}}$ of the top jet.
2. The N -subjettiness observables τ_{21} and τ_{32} of the top jet [55–59].

Those observables are shown to be important in identifying the color structure of the hard process [59, 67, 68]. In our scenario, the signal results from a color-singlet resonance, while the background comes from QCD process, and the jet mass and N -subjettiness can help to reveal this difference [67]. Moreover, such jet substructures can be more independent on resonance characteristics and kinematics.

Some distributions of these high-level observables are shown in Fig. 4. Note that the spin correlations as well as the jet substructure observables are rather insensitive to the width of ρ , as expected. For the 5 TeV resonance, the mass peak of $M_{t\bar{t}} \sim 5$ TeV almost disappears for $\Gamma_\rho/M_\rho \geq 10\%$; instead, there is a peak ~ 1 TeV, due to the parton-distribution support of off-shell effects and hard p_T cuts. Most identified top-jets in both signal and background originate correctly from the top quark, thus the differences shown in the distributions of $M_{j_{top}}$ and τ_{32} come from the color structure of the hard process. For example, the background’s $M_{j_{top}}$ distribution is slightly broader and the τ_{32} is slightly bigger than the signals. This is because the top-jets from QCD $t\bar{t}$ are color connected with the initial state, consequently having more radiations. Using these “all observables” (i.e. sum of low-

and high-level observables) as inputs, we train a new set of DNNs; best network configurations are again surveyed and detailed in Tables 3 and 4 of the Appendix.

We compare the performances of original and new DNNs using receiver operating characteristic (ROC) curves. The area under curve (AUC) is used as a metric of the performance. Some of the comparisons are shown in Fig. 5. First, in the resolved region as shown in the top panel, we found that there is only little change on ROC curves by adding high-level observables. Not only AUC, but also background efficiencies show small change. This means that the inclusion of high-level observables does not yield the improvement of accuracy; the original DNN had learned those high-level features successfully from low-level inputs.

In the boosted region, while the $M_{t\bar{t}}$, $M_{j_{top}}$ and spin correlations can be derived from the four momenta of reconstructed objects, the N -subjettiness cannot be inferred from the low-level inputs. Therefore, adding high-level features can bring improvements. As shown with ROC in the bottom two panels of Fig. 5, the improvement is sizable for $M1\Gamma4$, while, however, relatively small for $M5\Gamma4$. This may be because the event topology of $M5$ boosted cases becomes so simple that many features are more correlated.

4.2 Ranking input observables by importance

Which information has been used most usefully by DNN in distinguishing a broad resonance against continuum background? To answer this, we attempt to identify which connections between which neurons and layers are weighted most importantly. Following Ref. [69], we define the learning speed of the j -th hidden layer as

$$v^{(j)} = \left| \frac{\partial \mathcal{L}_{\text{loss}}(w, b)}{\partial \mathbf{b}^{(j)}} \right|, \quad (9)$$

where $\mathbf{b}^{(j)}$ is the bias vector of the j -th hidden layer, while $\mathcal{L}_{\text{loss}}$ is the loss function. As the target of machine learning is to find the global minimum of $\mathcal{L}_{\text{loss}}$, the $v^{(j)}$ approximately reflects the training sensitivity of a specific layer. When training the DNN, the larger $v^{(j)}$ a layer acquires, the more important it is. We found that for all individual benchmark cases

(a) Examples of distributions in the resolved region.

(b) Examples of distributions in the boosted region.

Fig. 4 Distributions of high-level observables in the **a** resolved and **b** boosted regions. We use these to train a new set of DNNs to test whether such high-level features were learned

$M_{t\bar{t}}$, the first hidden layer has the highest learning speed several times larger than that of other layers. For example, for $M1\Gamma4$ case in the resolved region, the learning speed is $v^{(1)} = 0.457$, $v^{(2)} = 0.086$, $v^{(3)} = 0.033$, $v^{(4)} = 0.016$ and $v^{(5)} = 0.008$. This means that good features are typically learned most efficiently in the first hidden layer.

For our DNN architecture described in Eq. (5), the weights of the first hidden layer form a $N_{in} \times N_{node}$ matrix, whose

Fig. 5 ROC curves, comparing the performance with (“all”) and without (“low”) high-level observables used to train DNNs. The AUC of each curve is also shown inside parenthesis

element is denoted as $w_{mn}^{(1)}$ with $m = 1, \dots, N_{in}$ and $n = 1, \dots, N_{node}$. As all the input features are rescaled to have average 0 and standard deviation 1, the magnitude of the weight $w_{mn}^{(1)}$ reflects the correlation strength between the m -th input and the n -th neuron in the first hidden layer. Motivated by this, we further define

$$W_m = \mathcal{N} \sqrt{\sum_{n=1}^{N_{node}} \left(w_{mn}^{(1)}\right)^2}, \quad (10)$$

as a measure of the importance of the m -th input feature. The normalization $\mathcal{N}$ is such that

$$\sum_{m=1}^{N_{in}} W_m = 1. \quad (11)$$

Figure 6 shows the W_m 's of each input observable from the DNN trained using both low- and high-level observables. Above all, the $M_{t\bar{t}}$ – that we expected to be less useful for a broad resonance – is still one of the most important observables even when the resonance is broad. This is particularly true for a low-mass broad resonance in the resolved region (upper panel). In the case of a heavy-resonance in the boosted region (lower panel), its importance is relatively reduced, partly because some invariant-mass information has been used in the selection of the boosted region. In such cases, the top-jet mass and transverse momentum which are somewhat correlated with $M_{t\bar{t}}$ and width can significantly complement the search, as shown in the bottom panel. In addition, the invariant mass of the top-jet is another important input feature because it reflects the color flow difference between

Fig. 6 The weight of each observable in the first layer of DNN, W_m defined in Eq. (10). $M_\rho = 1$ TeV in the resolved region (upper panel) and $M_\rho = 5$ TeV in the boosted region (lower panel). High-level observ-

ables are also used in training in order to get those weights. For notations, E_T and ϕ^{E_T} in the figures stand for $\not{E}_T$ and $\phi^{\not{E}_T}$, respectively. $c_{tl,th}^C$ and $c_{tl,th}^{CS}$ are short for $\cos \theta_{tl,th}^{CS}$, and $c_{1,2}^M$ and $c_{1,2}^{Mus}$ for $\cos \theta_{1,2}^{Mus}$.

signals and background. On the other hand, N -subjettinesses again turn out to be relatively less useful.

Remarkably, there are much other useful information, particularly from angular distributions $\eta^{\ell,j}$ and $\cos \theta_{1,2}^{Mus}$. From Figs. 1 and 4, we can see that these observables are relatively uncorrelated with the resonance width. We have indeed checked that the cross entropies [70] between these observables and $M_{t\bar{t}}$, which can quantify their correlations, are not so high. As we will see in the next subsection, these information are useful even in the off-shell region away from the resonance, hence less correlated with the width. Thus, these features are useful in search of broad resonances. This may also imply that narrow-resonance searches can be improved by adding off-resonance information; this is partly because a large fraction of signals is still from low-energy off-resonance region where parton-luminosity support is much

larger (although buried under larger backgrounds). We leave this for a future study.

4.3 Planing away $M_{t\bar{t}}$

We have observed that $M_{t\bar{t}}$ is still important, but there are indeed uncorrelated useful information. How much is discovery capability attributed to those uncorrelated (whether known or unknown) information? Using the data planing method [29, 71], we plane away the feature in the invariant mass spectrum. We attach a weight to each event so that the weighted distribution of $M_{t\bar{t}}$ becomes flat for both signals and backgrounds; the details of chosen network configurations and more results are described in Table 5 of the Appendix. A new set of DNNs trained with such planed data must learn information uncorrelated with $M_{t\bar{t}}$, and the dif-

ference between the performance with/without $M_{t\bar{t}}$ offers a quantitative answer to the question “how much information it is beyond the invariant mass”.

In practice, to avoid large fluctuations, we use only $M_{t\bar{t}} \in [0.5, 3]$ TeV region with 20 GeV bin size for all signal cases. This means that for 5 TeV signals, we consider only off-resonance events; note that the majority of signal is from the low-energy region supported by larger parton luminosities.

After $M_{t\bar{t}}$ planed away, the classification accuracies reduce from $\geq 80\%$ to $\geq 73\%$ for $M_\rho = 1$ TeV in the resolved region and from $\geq 65\%$ to $\geq 62\%$ in the boosted region. For $M_\rho = 5$ TeV cases, accuracies reduce from $\geq 76\%$ to $\geq 63\%$ in the boosted region. As accuracies are still significantly higher than random guess (i.e. 50%), we conclude that DNNs still have some capabilities to distinguish signals from background, even though they are blind to $M_{t\bar{t}}$ and most events are from off-resonance region (for 5 TeV cases). Clearly, on top of $M_{t\bar{t}}$ and width, the original DNNs had learned extra information (such as aforementioned angular correlations).

Indeed, we have checked that the weights W_m for various angular and angular-correlation observables, after planing the $M_{t\bar{t}}$, are relatively high. From Figs. 1 and 4, one can also see that they are largely independent on the width. The helicity conservation (hence, angular correlations) can hold somewhat independently of the invariant mass, as the range of the invariant mass considered is always much larger than the top mass. Thus, we conclude that much of the angular information can be from off-resonance region, and such off-resonance information (although buried under larger backgrounds) can enhance discovery power. As a result, as shown in Fig. 3, final performance is not only improved but became rather insensitive to the resonance width.

A final remark is that there could still be *unknown* (to us) useful information that are not identified in our analysis.

5 Conclusion

We have found that, in an attempt to develop methods to discover broad $t\bar{t}$ resonances, $M_{t\bar{t}}$ is still one of the most important observables, but additional information from both on- and off-resonance regions can significantly enhance discovery capability. As a result, the cross section upper limits can be improved by $\sim 60\%$ for $\Gamma_\rho/M_\rho \sim 40\%$, and the improved LHC sensitivities do not strongly depend on the width of a resonance. As resonances in new physics beyond the SM are easily broad, our learnings and techniques can be used to efficiently search for them.

The most useful observables turn out to be $M_{t\bar{t}}$ (even for broad resonances), $p_T^{j_{\text{top}}}$, $M_{j_{\text{top}}}$, angular distributions and color correlations. The usefulness of $M_{t\bar{t}}$ even for broad-resonance searches is not necessarily obvious, a priori. But correlated observables such as $p_T^{j_{\text{top}}}$ are found to further complement. Angular information (some of whose contributions come from off-resonance region) and $M_{j_{\text{top}}}$ (which can measure color flow structures irrespective of resonance characteristics) are relatively uncorrelated with the width and $M_{t\bar{t}}$, making improved LHC sensitivities less dependent on the width. Lastly, as we trained using only low-level inputs, our results also show that high-level observables such as $M_{t\bar{t}}$ are effectively well learned by DNN.

We have assessed these machine-learned information in three ways: by explicitly testing those high-level observables, by ranking input (low and/or high) observables using weights of the network, and by planing away features correlated with $M_{t\bar{t}}$. Notably, after all, there can still be unknown useful information that are not easily identified in our analysis. Thus, being able to communicate more efficiently with networks will enable better explorations of the nature, beyond what we know.

Acknowledgements We would like to thank Shawn Jia, Jinmian Li, Hui Luo, Tao Xu, Daneng Yang and Zhao-Huan Yu for discussions and the anonymous referee for useful suggestions. SJ and KPX are supported by Grant Korea NRF 2015R1A4A1042542, NRF 2017R1D1A1B03030820, SJ also by POSCO Science Fellowship, and DL by NRF 0426-20170003, NRF 0409-20190120.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors' comment: The size of the training/test dataset is too large (~ 10 GB) to be shared online. The dataset can be obtained by sending email to us.]

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 SCOAP³.

A The chosen DNN configurations and their performances

The selected DNN configurations for $M_\rho = 1$ and 5 TeV are listed in Tables 3 and 4, respectively. The selection criteria are described in Sect. 3.3. The epochs when we cut the train-

Table 3 The selected networks for $M_\rho = 1$ TeV. N_{epoch} is the epoch number when we cut the training

Benchmark case	Kinematic region	Observables	$N_{\text{hidden}}, N_{\text{node}}, L_r, D_r, B_s, N_{\text{epoch}}$	Classification accuracy (%)
M1 Γ 1	Resolved	Low-level	5, 200, 0.001, 0.2, 10^3 , 150	85.2
		All	5, 200, 0.001, 0.2, 10^3 , 100	85.1
	Boosted	Low-level	5, 200, 0.001, 0.2, 10^4 , 55	67.9
		All	4, 300, 0.001, 0.3, 10^3 , 40	70.1
M1 Γ 2	Resolved	Low-level	4, 300, 0.003, 0.2, 10^3 , 35	83.2
		All	5, 200, 0.001, 0.2, 10^3 , 60	83.2
	Boosted	Low-level	5, 200, 0.001, 0.2, 10^4 , 45	65.8
		All	4, 300, 0.003, 0.3, 10^4 , 40	68.2
M1 Γ 3	Resolved	Low-level	4, 300, 0.001, 0.2, 10^3 , 30	81.6
		All	5, 200, 0.001, 0.2, 10^3 , 40	81.6
	Boosted	Low-level	4, 300, 0.003, 0.2, 10^4 , 30	65.1
		All	4, 300, 0.003, 0.3, 10^4 , 40	67.0
M1 Γ 4	Resolved	Low-level	5, 200, 0.001, 0.2, 10^3 , 80	80.8
		All	5, 200, 0.001, 0.2, 10^3 , 40	80.6
	Boosted	Low-level	4, 300, 0.001, 0.2, 10^4 , 20	64.3
		All	4, 300, 0.001, 0.3, 10^3 , 40	66.7

Table 4 The selected networks for $M_\rho = 5$ TeV. N_{epoch} is the epoch number when we cut the training

Benchmark case	Kinematic region	Observables	$N_{\text{hidden}}, N_{\text{node}}, L_r, D_r, B_s, N_{\text{epoch}}$	Classification accuracy (%)
M5 Γ 1	Boosted	Low-level	4, 300, 0.001, 0.2, 10^3 , 20	79.5
		All	5, 200, 0.001, 0.1, 10^4 , 30	80.5
M5 Γ 2	Boosted	Low-level	4, 300, 0.003, 0.2, 10^3 , 40	78.2
		All	5, 200, 0.001, 0.1, 10^4 , 45	79.1
M5 Γ 3	Boosted	Low-level	4, 300, 0.003, 0.2, 10^3 , 30	77.4
		All	4, 300, 0.003, 0.3, 10^4 , 40	78.4
M5 Γ 4	Boosted	Low-level	4, 300, 0.003, 0.2, 10^3 , 45	76.8
		All	5, 200, 0.001, 0.3, 10^4 , 45	77.8

Table 5 The accuracy reach of the chosen neural networks before and after planing away $M_{t\bar{t}}$. The configurations of the DNN's are listed in Tables 3 and 4

Models	M1 Γ 1		M1 Γ 2		M1 Γ 3		M1 Γ 4		M5 Γ 1	M5 Γ 2	M5 Γ 3	M5 Γ 4
Kinematic region	Resolved	Boosted	Resolved	Boosted	Resolved	Boosted	Resolved	Boosted	Boosted	Boosted	Boosted	Boosted
Low-level input (%)	85.2	67.9	83.2	65.8	81.6	65.1	80.8	64.3	79.5	78.2	77.4	76.8
Planing away $M_{t\bar{t}}$ (%)	76.8	63.7	75.3	62.7	74.1	62.1	73.0	62.3	65.3	65.8	63.7	64.1

ing are listed in the forth columns. One can see that for a individual signal benchmark in a given kinematic region, the DNN with low-level observables usually requires a longer training epoch than the DNN with all observables, if they have the same configurations. That is because the DNN needs more time to learn about the physics in the signal process, if no hint is given to it. The classification accuracies (on the validation/test data) of the networks are given in the fifth columns.

Table 5 shows the accuracy reach of the DNNs before and after planing away the key observable $M_{t\bar{t}}$. The data of the second row, i.e. the accuracies before planing, are taken from the fifth columns of Tables 3 and 4. While the accuracies after planing listed in the third row are obtained by the weighted training described in Sect. 4.3.

References

1. D. Barducci, A. Belyaev, S. De Curtis, S. Moretti, G.M. Pruna, Exploring Drell–Yan signals from the 4D Composite Higgs Model at the LHC. *JHEP* **04**, 152 (2013). [arXiv:1210.2927](#)
2. D. Greco, D. Liu, Hunting composite vector resonances at the LHC: naturalness facing data. *JHEP* **12**, 126 (2014). [arXiv:1410.2883](#)
3. D. Barducci, C. Delaunay, Bounding wide composite vector resonances at the LHC. *JHEP* **02**, 055 (2016). [arXiv:1511.01101](#)
4. D. Liu, L.-T. Wang, K.-P. Xie, Broad composite resonances and their signals at the LHC. [arXiv:1901.01674](#)
5. R. Kelley, L. Randall, B. Shuve, Early (and later) LHC search strategies for broad dimuon resonances. *JHEP* **02**, 014 (2011). [arXiv:1011.0728](#)
6. S. Ask, J.H. Collins, J.R. Forshaw, K. Joshi, A.D. Pilkington, Identifying the colour of TeV-scale resonances. *JHEP* **01**, 018 (2012). [arXiv:1108.2396](#)
7. ATLAS Collaboration, M. Aaboud et al., Search for heavy particles decaying into top-quark pairs using lepton-plus-jets events in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Eur. Phys. J. C* **78**(7), 565 (2018). [arXiv:1804.10823](#)
8. K.J.F. Gaemers, F. Hoogeveen, Higgs production and decay into heavy flavors with the gluon fusion mechanism. *Phys. Lett. B* **146**, 347–349 (1984)
9. D. Dicus, A. Stange, S. Willenbrock, Higgs decay to top quarks at hadron colliders. *Phys. Lett. B* **333**, 126–131 (1994). [arXiv:hep-ph/9404359](#)
10. N. Craig, F. D’Eramo, P. Draper, S. Thomas, H. Zhang, The hunt for the rest of the Higgs bosons. *JHEP* **06**, 137 (2015). [arXiv:1504.04630](#)
11. S. Jung, J. Song, Y.W. Yoon, Dip or nothingness of a Higgs resonance from the interference with a complex phase. *Phys. Rev. D* **92**(5), 055009 (2015). [arXiv:1505.00291](#)
12. S.Y. Choi, J. Kalinowski, Y. Liao, P.M. Zerwas, H/A Higgs mixing in CP-noninvariant supersymmetric theories. *Eur. Phys. J. C* **40**, 555–564 (2005). [arXiv:hep-ph/0407347](#)
13. J.R. Ellis, J.S. Lee, A. Pilaftsis, CERN LHC signatures of resonant CP violation in a minimal supersymmetric Higgs sector. *Phys. Rev. D* **70**, 075010 (2004). [arXiv:hep-ph/0404167](#)
14. M. Carena, Z. Liu, Challenges and opportunities for heavy scalar searches in the $t\bar{t}$ channel at the LHC. *JHEP* **11**, 159 (2016). [arXiv:1608.07282](#)
15. CMS Collaboration, A.M. Sirunyan et al., Search for resonant $t\bar{t}$ production in proton–proton collisions at $\sqrt{s} = 13$ TeV. *JHEP* (2018). [arXiv:1810.05905 \(submitted\)](#)
16. C.M.S. Collaboration, A.M. Sirunyan et al., Search for narrow and broad dijet resonances in proton–proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter mediators and other new particles. *JHEP* **08**, 130 (2018). [arXiv:1806.00843](#)
17. ATLAS Collaboration, M. Aaboud et al., Search for resonances in the mass distribution of jet pairs with one or two jets identified as b -jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Lett. B* **759**, 229–246 (2016). [arXiv:1603.08791](#)
18. ATLAS Collaboration, M. Aaboud et al., Search for new high-mass phenomena in the dilepton final state using 36 fb^{-1} of proton–proton collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector. *JHEP* **10**, 182 (2017). [arXiv:1707.02424](#)
19. D. Liu, L.-T. Wang, K.-P. Xie, Prospects of searching for composite resonances at the LHC and beyond. *JHEP* **01**, 157 (2019). [arXiv:1810.08954](#)
20. J.H. Collins, K. Howe, B. Nachman, Anomaly detection for resonant new physics with machine learning. *Phys. Rev. Lett.* **121**(24), 241803 (2018). [arXiv:1805.02664](#)
21. J.H. Collins, K. Howe, B. Nachman, Extending the search for new resonances with machine learning. *Phys. Rev. D* **99**(1), 014038 (2019). [arXiv:1902.02634](#)
22. J. Hajer, Y.-Y. Li, T. Liu, H. Wang, Novelty detection meets collider physics. [arXiv:1807.10261](#)
23. P. Baldi, P. Sadowski, D. Whiteson, Searching for exotic particles in high-energy physics with deep learning. *Nat. Commun.* **5**, 4308 (2014). [arXiv:1402.4735](#)
24. P. Baldi, K. Cranmer, T. Faucett, P. Sadowski, D. Whiteson, Parameterized neural networks for high-energy physics. *Eur. Phys. J. C* **76**(5), 235 (2016). [arXiv:1601.07913](#)
25. H. Luo, M.-X. Luo, K. Wang, T. Xu, G. Zhu, Quark jet versus gluon jet: deep neural networks with high-level features. [arXiv:1712.03634](#)
26. J. Lee, N. Chanon, A. Levin, J. Li, M. Lu, Q. Li, Y. Mao, Polarization fraction measurement in same sign WW scattering using deep learning. [arXiv:1812.07591](#)
27. J. Pearkes, W. Fedorko, A. Lister, C. Gay, Jet constituents for deep neural network based top quark tagging. [arXiv:1704.02124](#)
28. J. Guo, J. Li, T. Li, F. Xu, W. Zhang, Deep learning for R -parity violating supersymmetry searches at the LHC. *Phys. Rev. D* **98**(7), 076017 (2018). [arXiv:1805.10730](#)
29. L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, A. Schwartzman, Jet-images—deep learning edition. *JHEP* **07**, 069 (2016). [arXiv:1511.05190](#)
30. J. Cogan, M. Kagan, E. Strauss, A. Schwartzman, Jet-images: computer vision inspired techniques for jet tagging. *JHEP* **02**, 118 (2015). [arXiv:1407.5675](#)
31. G. Li, Z. Li, Y. Wang, Y. Wang, Improving the measurement of Higgs boson–gluon coupling using convolutional neural networks at e^+e^- colliders. [arXiv:1901.09391](#)
32. K. Fraser, M.D. Schwartz, Jet charge and machine learning. *JHEP* **10**, 093 (2018). [arXiv:1803.08066](#)
33. G. Louppe, K. Cho, C. Becot, K. Cranmer, QCD-aware recursive neural networks for jet physics. *JHEP* **01**, 057 (2019). [arXiv:1702.00748](#)
34. M. Abdughani, J. Ren, L. Wu, J.M. Yang, Probing stop with graph neural network at the LHC. [arXiv:1807.09088](#)
35. J. Ren, L. Wu, J.M. Yang, Unveiling CP property of top–Higgs coupling with graph neural networks at the LHC. [arXiv:1901.05627](#)
36. I. Henrion, J. Brehmer, J. Bruna, K. Cho, K. Cranmer, G. Louppe, G. Rochette, Neural message passing for jet physics
37. D. Guest, K. Cranmer, D. Whiteson, Deep learning and its application to LHC physics. *Annu. Rev. Nucl. Part. Sci.* **68**, 161–181 (2018). [arXiv:1806.11484](#)
38. M. Abdughani, J. Ren, L. Wu, J.M. Yang, J. Zhao, Supervised deep learning in high energy phenomenology: a mini review. [arXiv:1905.06047](#)
39. D. Liu, R. Mahbubani, Probing top–antitop resonances with $t\bar{t}$ scattering at LHC14. *JHEP* **04**, 116 (2016). [arXiv:1511.09452](#)
40. H. An, X. Ji, L.-T. Wang, Light dark matter and Z' dark force at colliders. *JHEP* **07**, 182 (2012). [arXiv:1202.2894](#)
41. A. Alloul, N.D. Christensen, C. Degrande, C. Duhr, B. Fuks, FeynRules 2.0—a complete toolbox for tree-level phenomenology. *Comput. Phys. Commun.* **185**, 2250–2300 (2014). [arXiv:1310.1921](#)
42. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.S. Shao, T. Stelzer, P. Torrielli, M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP* **07**, 079 (2014). [arXiv:1405.0301](#)
43. M. Czakon, A. Mitov, Top++: a program for the calculation of the top–pair cross-section at hadron colliders. *Comput. Phys. Commun.* **185**, 2930 (2014). [arXiv:1112.5675](#)

44. M. Czakon, P. Fiedler, A. Mitov, Total top-quark pair-production cross section at hadron colliders through $O(\alpha_s^4)$. Phys. Rev. Lett. **110**, 252004 (2013). [arXiv:1303.6254](#)
45. M. Czakon, A. Mitov, NNLO corrections to top pair production at hadron colliders: the quark–gluon reaction. JHEP **01**, 080 (2013). [arXiv:1210.6832](#)
46. M. Czakon, A. Mitov, NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels. JHEP **12**, 054 (2012). [arXiv:1207.0236](#)
47. P. Bärnreuther, M. Czakon, A. Mitov, Percent level precision physics at the Tevatron: first genuine NNLO QCD corrections to $q\bar{q} \rightarrow t\bar{t} + X$. Phys. Rev. Lett. **109**, 132001 (2012). [arXiv:1204.5201](#)
48. M. Cacciari, M. Czakon, M. Mangano, A. Mitov, P. Nason, Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation. Phys. Lett. B **710**, 612–622 (2012). [arXiv:1111.5869](#)
49. T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C.O. Rasmussen, P.Z. Skands, An introduction to PYTHIA 8.2. Comput. Phys. Commun. **191**, 159–177 (2015). [arXiv:1410.3012](#)
50. DELPHES 3 Collaboration, J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi, DELPHES 3, A modular framework for fast simulation of a generic collider experiment. JHEP **02**, 057 (2014). [arXiv:1307.6346](#)
51. ATLAS Collaboration, M. Aaboud et al., Search for new phenomena in events with same-charge leptons and b -jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. [arXiv:1807.11883](#)
52. Performance of jet substructure techniques in early $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector, Tech. Rep. ATLAS-CONF-2015-035. CERN, Geneva (2015)
53. Boosted hadronic top identification at ATLAS for early 13 TeV data, Tech. Rep. ATL-PHYS-PUB-2015-053. CERN, Geneva (2015)
54. J. Thaler, L.-T. Wang, Strategies to identify boosted tops. JHEP **07**, 092 (2008). [arXiv:0806.0023](#)
55. D.E. Kaplan, K. Rehermann, M.D. Schwartz, B. Tweedie, Top tagging: a method for identifying boosted hadronically decaying top quarks. Phys. Rev. Lett. **101**, 142001 (2008). [arXiv:0806.0848](#)
56. J. Thaler, K. Van Tilburg, Identifying boosted objects with N-subjettiness. JHEP **03**, 015 (2011). [arXiv:1011.2268](#)
57. J. Thaler, K. Van Tilburg, Maximizing boosted top identification by minimizing N-subjettiness. JHEP **02**, 093 (2012). [arXiv:1108.2701](#)
58. T. Plehn, M. Spannowsky, Top tagging. J. Phys. **G39**, 083001 (2012). [arXiv:1112.4441](#)
59. G. Kasieczka, T. Plehn, T. Schell, T. Streblar, G.P. Salam, Resonance searches with an updated top tagger. JHEP **06**, 203 (2015). [arXiv:1503.05921](#)
60. F. Chollet et al., Keras. <https://keras.io> (2015)
61. M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., Tensorflow: a system for large-scale machine learning. In: *OSDI*, vol. 16, pp. 265–283 (2016)
62. D. Yang, Q. Li, Probing the dark sector through mono-Z boson leptonic decays. JHEP **02**, 090 (2018). [arXiv:1711.09845](#)
63. **The ATLAS Collaboration, The CMS Collaboration, The LHC Higgs Combination Group** Collaboration, Procedure for the LHC Higgs boson search combination in Summer 2011, Tech. Rep. CMS-NOTE-2011-005. ATL-PHYS-PUB-2011-11. CERN, Geneva (2011)
64. K. Datta, A. Larkoski, How much information is in a jet? JHEP **06**, 073 (2017). [arXiv:1704.08249](#)
65. J.C. Collins, D.E. Soper, Angular distribution of dileptons in high-energy hadron collisions. Phys. Rev. D **16**, 2219 (1977)
66. E. Richter-Was, Z. Was, Separating electroweak and strong interactions in Drell-Yan processes at LHC: leptons angular distributions and reference frames. Eur. Phys. J. C **76**(8), 473 (2016). [arXiv:1605.05450](#)
67. K. Joshi, A.D. Pilkington, M. Spannowsky, The dependency of boosted tagging algorithms on the event colour structure. Phys. Rev. D **86**, 114016 (2012). [arXiv:1207.6066](#)
68. D.E. Soper, M. Spannowsky, Finding physics signals with event deconstruction. Phys. Rev. D **89**(9), 094005 (2014). [arXiv:1402.1189](#)
69. M.A. Nielsen, *Neural networks and deep learning*, vol. 25. Determination Press, USA (2015)
70. T. Roxlo, M. Reece, Opening the black box of neural nets: case studies in stop/top discrimination. [arXiv:1804.09278](#)
71. S. Chang, T. Cohen, B. Ostidek, What is the machine learning? Phys. Rev. D **97**(5), 056009 (2018). [arXiv:1709.10106](#)