

Improving the direct determination of $|V_{ts}|$ using deep learning

Jeewon Heo¹, Woojin Jang¹, Jason S. H. Lee^{1,*}, Youn Jung Roh¹, Ian James Watson^{1,†}, and Seungjin Yang²

¹Department of Physics, University of Seoul, Seoul 02504, Republic of Korea

²Department of Physics, Kyung Hee University, Seoul 02453, Republic of Korea



(Received 16 February 2025; accepted 11 August 2025; published 11 September 2025)

An s -jet tagging approach to determine the Cabibbo-Kobayashi-Maskawa matrix component $|V_{ts}|$ directly in the dileptonic final state events of the top pair production in proton-proton collisions has been previously studied by measuring the branching fraction of the decay of one of the top quarks by $t \rightarrow sW$. The main challenge is improving the discrimination performance between strange jets from top decays and other jets. This study proposes novel jet discriminators, called DiSaJa, using a Transformer-based deep learning method. The first model, DiSaJa-H, utilizes multidomain inputs (jets, leptons, and missing transverse momentum). An additional model, DiSaJa-L, further improves the setup by using lower-level jet constituent information, rather than the high-level clustered information. DiSaJa-L is a novel model that combines low-level jet constituent analysis with event classification using multidomain inputs. The model performance is evaluated via a CMS-like LHC Run 2 fast simulation by comparing various statistical test results to those from a Transformer-based jet classifier which considers only the individual jets. This study shows that the DiSaJa models have significant performance gains over the individual jet classifier, and we show the potential of the measurement during Run 3 of the LHC and the High-Luminosity LHC.

DOI: [10.1103/612r-ycj7](https://doi.org/10.1103/612r-ycj7)

I. INTRODUCTION

The Cabibbo-Kobayashi-Maskawa (CKM) matrix is the 3×3 unitary complex matrix that gives the strength of the charged-current weak interaction between the quark generations in the Standard Model (SM) [1]. A global fit has been performed to constrain its components using measurements of various aspects of the CKM matrix and by imposing the SM condition of unitarity [2]. Although the fit gives precise values for each CKM component, further measurements are necessary to test the validity of the unitarity condition. In particular, the unitarity is no longer valid in several beyond the SM (BSM) theories [3]. Therefore, direct measurement of the components should be performed to test the SM consistency and constrain BSM scenarios.

In this paper, we focus on the measurement potential of the third-row component $|V_{ts}|$, whose squared value gives the branching ratio of the decay of the top quark to the strange quark and a W boson in the SM. In the global fit of the CKM under the SM conditions, the value of $|V_{ts}|$ is

$4.110^{+0.083}_{-0.072} \times 10^{-2}$ [2]. There have been several studies for measuring the component indirectly, which are used in the global fit. For example, $|V_{ts}|$ is determined indirectly using the $B_s^0 - \bar{B}_s^0$ oscillation frequency [4–6] and decay constant parameters from lattice QCD results [7], which results in $|V_{ts}| = 4.15 \pm 0.09 \times 10^{-2}$ [2]. However, as the indirect measurements rely on loop processes, there could be BSM contributions and therefore these measurements could yield results that differ from the true value of $|V_{ts}|$. For example, BSM models with additional quark generations allow $|V_{ts}|$ to be as large as 0.2 [3].

There are several measurements for the model-independent direct determination of the V_{tx} components, where x is d , s , and b . For instance, there are recent analyses with the ATLAS and CMS detectors using 13 TeV data where the single top process probes the tWq vertices in production and decay in the t -channel. The CMS study gives limits of $|V_{ts}|^2 + |V_{td}|^2 < 0.057$ and $|V_{ts}|^2 + |V_{td}|^2 < 0.06$ at the 95% confidence level (CL) under SM CKM unitarity and after relaxing the SM constraint, respectively [8]. The ATLAS study uses the fact that the d -quark is a valence quark of the proton while the s -quark is a sea quark to give the separate limits $f_{LV}|V_{ts}| < 0.58$ and $f_{LV}|V_{td}| < 0.23$ at the 95% CL, where f_{LV} is the left-handed form factor [9]. Additionally, previous studies have proposed the direct determination using a light-flavor jet tagging approach to discriminate strange jets from the $t \rightarrow sW$ decay in the top pair production process for $|V_{ts}|$ [10,11] or the b -jets from $t \rightarrow bW$ for $|V_{tb}|$ [12].

*Contact author: jason.lee@uos.ac.kr

†Contact author: ian.james.watson@cern.ch

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

In this study, we expand on the jet tagging strategy for measuring $|V_{ts}|$ using a deep learning (DL) approach. The direct s -tagging approach is challenging due to the lack of statistics for signal events from $t \rightarrow sW$ compared to $t \rightarrow bW$, which is the most dominant background process, as the ratio of the signal to the background decay is given by $\frac{|V_{tb}|^2}{|V_{ts}|^2} \simeq 590$. Consequently, improving the separation power between the signal and background jets is crucial.

There are several previous studies of s -tagging jets in collider physics. At the LEP e^+e^- collider, the DELPHI Collaboration studied the strange-quark forward-backward asymmetry around the Z^0 peak, which used high-energy charged kaons to tag the strange quark [13]. The SLD at the SLC used high-energy charged kaons and K_S^0 to tag $Z^0 \rightarrow s\bar{s}$ events and measured the parity-violating coupling of the Z to s -quark [14]. A future e^+e^- collider study focuses on tagging strange-quark pairs from the decay of a Higgs boson $H \rightarrow s\bar{s}$, presenting an s -jet discrimination variable produced by tagging the charged kaons and K_S^0 inside of each jet [15]. For the proposed International Linear Collider, a study has been performed to show the measurement prospects for observing the rare decay $H \rightarrow s\bar{s}$ [16]. The study uses information from the jet constituents combined with high-level jet variables in order to discriminate the s -jets from the Higgs decay from background jets. There are also various proposals for s -jet tagging in hadron colliders, for example using track information with a Long Short-term Memory (LSTM) recurrent neural network to discriminate s -quarks from other light quarks [17], and using also calorimeter and Cherenkov detector information [18]. Other studies propose using jet images passed into a convolutional neural network [19], or using the jet particle information passed through a ParticleNet graph neural network [20]. Transformer-based jet tagging in future e^+e^- colliders has also been investigated and applied to the study of $Z^0 \rightarrow s\bar{s}$ [21]. Another study investigates the performance of a graph-attention network or particle transformer to tag strange jets with the High-Luminosity LHC (HL-LHC) [22]. In these previous studies, the s -jet is tagged and then utilized for further analysis.

In contrast, we propose a novel method to separate strange jets originating from top decays, using the full event information, and starting from a self-attention-based network, SaJa [23]. SaJa was originally developed for the assignment of jets to partons in the $t\bar{t}$ all-hadronic channel, where large QCD multijet backgrounds dominate the analysis. We extend the SaJa model to apply to the dilepton channel events of top pair production, and we call these new models DiSaJa. Using the dileptonic channel, there are fewer background jets in each event, due to the reduced jet activity in an event compared to the other top pair decay channels. To reflect the diverse decay products in the dilepton channel, DiSaJa-H employs dedicated embedding networks for leptons, jets, and missing energy to process all the reconstructed physics objects in an event.

Numerous studies [24–27] have reported that DL models using jet constituents as inputs demonstrate outstanding performance in object-level tasks such as flavor tagging. However, for event-level tasks, such as signal-background discrimination and jet-parton assignment, the representation of input jets has been restricted to the format of high-level feature variables rather than their constituents [23,28,29]. In this study, we produce an additional model, DiSaJa-L, which replaces the jet embedding layer in DiSaJa-H with a dedicated embedding network that processes jet constituents as inputs, enabling the model to learn jet representations optimized for this analysis. DiSaJa-L is thus a new general-purpose model that incorporates both low-level jet constituent analysis and multidomain inputs able to process the complete information available in hadron collider data.

This paper is organized as follows: Sec. II describes the event generation and detector simulation used in our analysis. In Sec. III, we present the object and event selection criteria. In Sec. IV, we explain the expanded SaJa networks, which we refer to as DiSaJa, applied to $t\bar{t}$ dilepton events and introduce a baseline model to evaluate the performance of the DiSaJa models. In Sec. V, we compare the model performance for the $|V_{ts}|$ measurement between two DiSaJa models and the baseline model with the simulated dataset. We also check the sensitivity of the measurement expected from the Run 3 and HL-LHC experiments [30] for evaluating the prospect of analyses performed at other integrated luminosities.

II. SIMULATION SETTINGS

We generate $t\bar{t}$ dilepton channel events with up to two additional partons in pp collisions at $\sqrt{s} = 13$ TeV at next-to-leading order (NLO) in QCD using MadGraph5_aMC@NLO2.6.5 [31] with NNPDF 3.1 [32]. The signal process is $t\bar{t}$ where one t -quark decays to an s -quark ($t\bar{t} \rightarrow sWbW$) while the background is $t\bar{t}$, where both t -quarks decay to b -quarks ($t\bar{t} \rightarrow bWbW$) and about 100M events are generated for each process. The inclusive $t\bar{t}$ cross section for $\sqrt{s} = 13$ TeV is calculated to be 831.76 pb, which is obtained at the next-to-next-to-leading-order (NNLO) QCD and next-to-next-to-leading-logarithmic soft-gluon resummation with Top++ [33]. We take the cross sections of dileptonic $t\bar{t} \rightarrow sWbW$ and $t\bar{t} \rightarrow bWbW$ to be 0.337 and 88.99 pb, respectively, by using the values of the inclusive cross section, the branching ratio of $W \rightarrow \ell\nu$ ($\ell = e, \mu, \tau$), V_{ts} , and V_{tb} [2]. While the most dominant background is the $t\bar{t} \rightarrow bWbW$, there are also non-negligible backgrounds from non- $t\bar{t}$ processes such as single top (ST), Drell-Yan (DY), and diboson (VV) production. The ST t -channel and tW -associated processes with no additional partons and DY events with two additional partons in the final state (DY + jj) are generated at the NLO accuracy in QCD using MadGraph5_aMC@NLO2.6.5

[31] with NNPDF 3.1 [32] and the number of generated events is about 220M, 220M, and 250M, respectively. In the ST generation, the W boson is forced to decay leptonically for more efficient event generation. For the generation of the DY process, the invariant mass of the final state lepton pair is set to be greater than 50 GeV. About 20M events for each VV process (WW , WZ , and ZZ) are generated using Pythia 8.240 [34] and their cross sections are set to 118.7 [35], 49.98 [36], and 16.91 pb [37], respectively, based on the NNLO QCD calculations.

After the matrix-element-level event generation, parton showering and hadronization are simulated with Pythia 8. The matrix-element events are jet-matched with the parton shower using the FxFx scheme [38] and the CP5 tuning parameters are used for modeling the underlying event [39]. After the simulation, the cross sections of the ST t -channel, the tW associated, and the DY are obtained as 73.45, 3.289, and 359.1 pb, respectively.

We use Delphes 3.4.2 [40] to simulate the response of a CMS-like detector. Delphes takes outputs from Pythia 8 and emulates the propagation of the particles in the magnetic field and the response of the particles in the detector's tracker and calorimeters. Using this information, Delphes produces reconstructed charged particle tracks (tracks) and neutral particles' energy depositions in the calorimeters (towers). These objects are used for reconstructing high-level objects, which are the isolated leptons and the jets made from clustering the tracks and towers. The kinematics of the tracks and towers objects are also summed and the transverse component of the result is negated to produce the missing transverse momentum ($\vec{p}_T^{\text{miss}}$ or MET) of the event. We change the default Delphes CMS card to reflect the Run 2 conditions by using update values for the ΔR for the lepton isolation, the jet clustering radius, and the b -tagging efficiency. The ΔR for lepton isolation is set to 0.3 (0.4) for electrons (muons). The anti- k_T algorithm is used for jet clustering with the jet radius $R = 0.4$ using Fastjet 3.3.2 [41], and the b -tagging efficiency is updated based on the efficiency distribution used by the CMS

experiment [42,43]. To emulate tracks in the CMS tracker, the track impact parameter is smeared and the resolution of the track transverse momentum is applied based on a CMS tracker performance study [44].

III. EVENT SELECTION

This analysis is performed with top pair production in the dilepton ee , $e\mu$, and $\mu\mu$ channels. We identify $t\bar{t}$ dilepton events using the standard selection criteria found in various CMS top analyses [45–47]. Charged leptons are selected using a cone-based relative isolation I_{rel} [40], and kinematic requirements. For muons, the isolation is required to be $I_{\text{rel}} < 0.15$ while electrons are selected when $I_{\text{rel}} < 0.0588$ (0.0571) in the barrel (end cap) region. Both flavors of lepton are required to be within $|\eta| < 2.4$, but electrons in the electromagnetic calorimeter transition gap region $1.44 < |\eta| < 1.57$ are excluded [48]. We select jets with $p_T > 30$ GeV and $|\eta| < 2.4$, vetoing jets where the distance from a selected lepton ΔR is less than 0.4. Among the remaining selected jets, jets are b -tagged according to the CMS b -tagging efficiency parametrized as a function of p_T .

We select events with exactly one lepton pair with opposite charges where the invariant mass $M_{\ell\ell}$ of the lepton pair is required to be greater than 20 GeV, and the p_T of the leading (subleading) lepton is required to be greater than 25 (20) GeV. For the same-flavor (SF) channel, we require $|M_{ll} - M_Z| > 15$ GeV, where the mass of the Z boson $M_Z \simeq 91$ GeV [2], to veto the Z boson background. Additionally, for the SF channel, we require that the missing transverse momentum $p_T^{\text{miss}} > 40$ GeV. We use events with at least two selected jets, where at most one jet is b -tagged.

We refer to jets originating from the parton q in top-quark decays $t \rightarrow qW$ as primary jets. Consequently, the signal jet is called the primary s -jet in this paper. Primary jets are identified as reconstructed jets matched to generator-level quarks from top-quark decays. Matching is performed by requiring that the distance ΔR between the




FIG. 1. Normalized distributions of the transverse momentum, pseudorapidity, and mass of reconstructed jets matched to prompt s - and b -quarks from top-quark decays in fully matched $t\bar{t} \rightarrow sWbW$ and $t\bar{t} \rightarrow bWbW$ events after the event selection.

parton and the jet satisfies $\Delta R < 0.4$. If there exist multiple ΔR -matched jets, which occur in less than 1% of signal events, the jet having the highest p_T is identified as the primary jet. In 85% of signal events, there is a jet matched to the $t \rightarrow sW$ parton.

Figure 1 presents the normalized distributions of transverse momentum, pseudorapidity, and mass for reconstructed jets that matched to partons from top-quark decays in fully matched $t\bar{t} \rightarrow sWbW$ and $t\bar{t} \rightarrow bWbW$ events, after the event selection. The near-identical shapes of these distributions reflect the treatment of both s - and b -quarks as massless at the matrix-element level. Small differences in p_T and mass between primary s - and b -jets stem from different hadronization processes and masses of b - and s -quarks. Moreover, primary b -jets in $t\bar{t} \rightarrow bWbW$ events are slightly less energetic than signal b -jets. This difference arises because the b -tagging efficiency increases with jet p_T , and the event selection rejects any event containing two or more b -tagged jets. As a result, high- p_T b -jets in the background are preferentially rejected, leading to an overall softer primary b -jet distribution in $t\bar{t} \rightarrow bWbW$ events.

IV. MACHINE LEARNING

The original saJa model, illustrated in Fig. 2, is designed for the task of jet assignment in fully hadronic top pair production and is built upon the Transformer encoder architecture [49]. It processes high-level jet features using a combination of feedforward networks (FFNs) and a

Transformer encoder block, which is displayed in Fig. 3. The FFN block consists of two layers of affine transformations, each followed by a Gaussian error linear unit activation function [50], with dropout applied to prevent overfitting [51]. The array of jet vectors is passed through the jet-wise FFN blocks. The encoder block, detailed below, allows for the interaction between the jet arrays through the use of the self-attention mechanism.

Generically, attention is a function that takes a source $\mathbf{S} \in \mathbb{R}^{M \times D_S}$ and a target $\mathbf{T} \in \mathbb{R}^{N \times D_T}$ as input and produces an output array of the same length as the target, where M and N are the lengths of arrays and each element in $\mathbf{S}$ ($\mathbf{T}$) is a vector of dimension D_S (D_T). The purpose of attention is to transform $\mathbf{T}$ into a rich contextual representation by extracting and integrating relevant information from $\mathbf{S}$; we say that $\mathbf{T}$ attends to $\mathbf{S}$. First, $\mathbf{T}$ is projected into $\mathbf{Q} \in \mathbb{R}^{N \times D_K}$, and $\mathbf{S}$ is projected into $\mathbf{K} \in \mathbb{R}^{M \times D_K}$ and $\mathbf{V} \in \mathbb{R}^{M \times D_V}$ using separate affine transformations. $\mathbf{Q}$, $\mathbf{K}$, and $\mathbf{V}$ are then passed through a scaled-dot-product attention function:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{D_K}}\right)\mathbf{V} \in \mathbb{R}^{N \times D_V}, \quad (1)$$

where softmax is applied to each row of the output of scaled-dot-product attention. Self-attention is a special case of attention where $\mathbf{S} = \mathbf{T}$. That is, a single set of objects attends to itself.

In the encoder block, multihead self-attention (MHSA) is used, which is a concatenation of N_{head} copies of the




FIG. 2. Architecture of the original saJa model (left) and the DiSaJa-H (right). N_x , D_x , and C denote the number of the object $x \in \{\text{jets, leptons, MET}\}$, the dimension size of the object x , and the number of output categories, respectively.




FIG. 3. Detailed network structure of blocks of feedforward network, encoder, and decoder. The decoder processes the output of the jet embedder as the target input and the output of the event encoder block as the source input.

scaled-dot-product attention described above. The encoder is comprised of N_{block} encoder blocks running sequentially, where each block consists of an MHSA block followed by an FFN block. The output of these blocks is added residually to the input arrays.

Unlike the original SaJa, the DiSaJa-H is designed to process multidomain inputs to utilize all of the objects in the dilepton final state events effectively. Figure 2

presents an overview of the DiSaJa-H architecture, including input embedding networks (or embedders), event encoder, decoder, and jet-wise classification head blocks. In the initial step, input features listed in Table I for each object (reconstructed jet, lepton, and MET) are embedded into the same dimensional space through each FFN block. The momentum components of the jet, the number of particles in the jet (for each category of particle), the jet energy

TABLE I. Features used as inputs in the models for each object type (jet, lepton, and MET).

Object	Variable	Definition
Jet	$p_T(j), \eta(j), \phi(j), M(j)$	Momentum components of jet
	N_{h^0}	Neutral hadron multiplicity
	$N_{h^\pm}$	Charged hadron multiplicity
	N_e	Electron multiplicity
	N_μ	Muon multiplicity
	N_p	Photon multiplicity
	$p_T D$	Jet energy sharing
	Jet axes	Lengths of ellipse
	Jet b -tag	Boolean indicating whether a jet is b -tagged or not
	Jet charge	Jet charge
Lepton	$p_T(\ell), \eta(\ell), \phi(\ell), M(\ell)$	Momentum components of lepton
	Lepton flavor	0 for e, 1 for μ
	Q_ℓ	Lepton charge
MET	$p_T^{\text{miss}}, \phi(p_T^{\text{miss}})$	Magnitude and azimuth angle of $\vec{p}_T^{\text{miss}}$




FIG. 4. Architecture of the jet constituent encoder, which can replace the jet high-level feature encoder. Track and tower features are fed into encoders and the jet constituent encoder learns jet representation.

sharing ($p_T D = \frac{\sqrt{\sum_i p_{T,i}^2}}{\sum_i p_{T,i}}$, where i indexes over particles inside jet) [52,53], the jet shape, the jet b -tagging information, and the jet charge ($Q_\kappa = \sum_{h \in \text{jet}} z_h^\kappa Q_h$, where $z_h = p_{T,h}/p_{T,\text{jet}}$, $\kappa = 0.3$) [54–56] are used as inputs to the jet array. For leptons, the momentum components, flavor, and charge are used. For the missing transverse momentum, its magnitude and azimuth angle are used as inputs. All input features are scaled to a range between 0 and 1 using min-max scaling. In DiSaJa-H, the separate object embedders allow different objects to be concatenated in a sequence by projecting them into the same dimensional space and they are then processed together by the event encoder, as in the original SaJa model. The encoder is followed by the decoder, which uses the Transformer

decoder architecture, and which takes the jet embedder's output as the target input and the event encoder's output as the source input. The MHSA block first processes the jet embedding input, and the output is combined with the event encoder output using multihead attention to integrate the full event information into each jet vector. The output of the decoder is passed to the jet-wise classification head, which assigns categorical scores for jets in each event using the final FFN head.

The DiSaJa-L model starts from the DiSaJa-H model as a base, and is augmented by utilizing arrays of jet constituent information as inputs instead of the high-level jet variables produced after the jet clustering. The low-level jet embedder is thus a drop-in replacement of the high-level jet embedder. Figure 4 shows the architecture of the low-level jet embedder, which is designed to extract the informative representations of jets from their constituents, which are the tracks and towers produced by Delphes. The input feature variables of tracks and towers are summarized in Table II. To address these differences, the low-level jet embedder includes two separate FFNs, which project tracks and towers into the same dimension, and are called the track and tower embedders, respectively. The outputs of track and tower embedders are then concatenated and passed into a jet constituent encoder, which also uses the Transformer encoder architecture described above. Then, the aggregate block averages the output of the jet constituent encoder over the constituent axis to produce a single vector per jet, which is used in the rest of the model, as in DiSaJa-H.

For the training and validation, we use a selected subsample of the generated $t\bar{t} \rightarrow sWbW$ and $t\bar{t} \rightarrow bWbW$ events. For training, we use around 1.1M events, which are required to contain a $t \rightarrow sW$ jet-parton matched jet, with no requirement of a $t \rightarrow bW$ matched jet. In the case of the background sample, we use about 0.5M events where both $t \rightarrow bW$ partons have jet matches and 0.4M events of unmatched events, where at least one of the $t \rightarrow bW$ jet-parton matches is missing. For model selection, we use around 275K signal events and 221K (121K of matched and 100K of unmatched) background events, passing the

TABLE II. Input features of jet constituents.

Variable	Definition
$p_T(P), \eta(P), \phi(P)$	Momentum components of particle
$\Delta\eta$	Difference of pseudorapidity between particle and jet axis
$\Delta\phi$	Difference of azimuthal angle ϕ between particle and jet axis
$\frac{p_T(P)}{p_T(j)}$	p_T of a constituent relative to jet p_T
p_T^{rel}	Particle momentum perpendicular to the jet axis
p_z^{rel}	Particle momentum in the direction of the jet axis
d_0	Transverse track impact parameter value
d_z	Longitudinal track impact parameter value
Q_P	Charge of particle
$E_{\text{EM}}, E_{\text{had}}$	Electromagnetic, hadronic energy in calorimeter

same matching requirements but chosen separately from the training samples, as the validation dataset.

The models are trained to classify jets into three groups: $t \rightarrow sW$, $t \rightarrow bW$, and other jets, which represent the jet categories of interest in the signal process $t\bar{t} \rightarrow sWbW$. The models are provided with the true labels from the simulation for the jet category of each jet in an event during training. We use jet-wise cross entropy as the objective function L for training, which is defined for each event as follows:

$$L(\theta) = \frac{1}{N} \sum_{j=1}^N \left(- \sum_{c \in \mathbb{C}} y_c^{(j)} \log \hat{y}_c^{(j)} \right) \quad (2)$$

where θ denotes the adjustable parameters of a model, N is the number of jets in the event, j indexes over the jets in the events, c indexes over the jet categories $\mathbb{C} = \{t \rightarrow sW, t \rightarrow bW, \text{other}\}$, $y_c^{(j)}$ is 1 for the true jet category c for jet j and 0 otherwise, and $\hat{y}_c^{(j)} = \hat{y}_c^{(j)}(\theta)$ is the model output for the jet j in category c . Model optimization is performed with the AdamW optimizer [57] with $\beta_1 = 0.9$, $\beta_2 = 0.999$, a learning rate of 0.0003, and a weight decay coefficient of 0.01. We use minibatch training, where each batch consists of 128 randomly sampled events for each training iteration. To handle variable-length jets and their constituents, input variables are zero padded to match the maximum length of inputs within a batch, allowing us to use all the jets in each event without truncation. We evaluate the loss on the validation set during training and hyperparameter optimization and select the model with the lowest loss.

While we train models to classify all the jets of the $t\bar{t} \rightarrow sWbW$ signal process, DiSaJa's output scores should also effectively discriminate signal events against background events. Further, since the difference between the $t\bar{t} \rightarrow sWbW$ signal process and the main background $t\bar{t} \rightarrow bWbW$ is the presence of the $t \rightarrow sW$ jet, we use the highest $t \rightarrow sW$ score within each event as a signal-background discriminant and the corresponding jet is referred to as the predicted primary s -jet. However, we found that models trained on only signal events, while achieving good assignment performance showed limited discrimination power between signal and background events. This challenge arises because the $t\bar{t} \rightarrow bWbW$ background process shares the same event topology with the $t\bar{t} \rightarrow sWbW$ signal, and other background processes can also mimic it. Because the $t\bar{t} \rightarrow bWbW$ background is statistically dominant after the final event selection, we explore incorporating $t\bar{t} \rightarrow bWbW$ background events into the training set. The impact of these different training configurations is evaluated by comparing the significance of excluding $|V_{ts}| = 0$, assuming an integrated luminosity of 138 fb^{-1} with only MC statistical uncertainty. The precise definition of the significance is given in Sec. V.

First, we train a DiSaJa-H model using the signal-only training set, which contains only jet-parton matched

$t\bar{t} \rightarrow sWbW$ events and constructs a baseline for the different training set configurations. The signal-only training set model achieves a significance of 2.60σ . The second configuration is based on a training set comprising both matched $t\bar{t} \rightarrow sWbW$ signal events and matched $t\bar{t} \rightarrow bWbW$ background events. Jets in matched $t\bar{t} \rightarrow bWbW$ events are labeled using jet-parton matching information as for $t\bar{t} \rightarrow sWbW$ events. The result with this configuration yields a significance of 3.24σ , demonstrating improved discrimination between signal and background. The training set for the final configuration also includes the unmatched $t\bar{t} \rightarrow bWbW$, and for these events all jets are labeled as other jets. This approach results in a significance of 4.45σ , the highest among the tested configurations. Based on these results, we use the final training configuration for the results of this study.

We optimize the hyperparameters of the DiSaJa-H model using the tree-structured Parzen estimator algorithm [58] within the Optuna framework [59]. The same hyperparameters are applied to DiSaJa-L. The hyperparameters for the model are D_{FFN} , N_{block} , N_{head} , and D_{model} , where D_{model} is $N_{\text{head}} \times$ dimension of each attention in MHSA, and are determined to be 1024, 2, 12, and 384, respectively.

For comparison to current DL methods, we use a baseline model which is a jet tagger based on the jet encoder section of DiSaJa-L, but passed into an s -tagging FFN rather than being used in DiSaJa. The input of the baseline classifier is an array of jet constituents from a single jet and uses the same features of the constituents, as listed in Table II, that are used for input into the jet encoder. It is trained to classify each jet into three categories: $t \rightarrow sW$, $t \rightarrow bW$, and other jets by minimizing the cross entropy of the model predictions. While the DiSaJa models process all jets and other objects simultaneously using the attention mechanism, providing outputs for all jets in an event at once, the baseline model processes jets individually, without considering their relationships with other jets. This gives a DL s -tagger approach for our baseline, using the per-jet low-level jet constituent information as in previous flavor-tagging-type studies.

V. RESULTS

We evaluate the performance of the DiSaJa models by comparing with the baseline model. The highest $t \rightarrow sW$ score within each event is used as the discriminant to distinguish between signal and background processes, and the jet with the highest $t \rightarrow sW$ score is referred to as the predicted primary s -jet. Figure 5 shows the distribution of the highest $t \rightarrow sW$ score in the signal sample. Events are categorized into three labels: correct (wrong), where a predicted primary s -jet is (is not) the genuine primary s -jet, and unmatched, representing events where the jet-parton matching fails. Figure 6 displays the fraction of the $t \rightarrow sW$ signal which is correctly assigned, wrongly assigned, and unmatched, with respect to the signal efficiency defined as




FIG. 5. Distribution of the highest $t \rightarrow sW$ score used as a discriminant in the signal sample, showing jets that are correctly assigned, wrongly assigned, and unmatched with partons. The ratios for these three categories are reflected in the distributions.




FIG. 6. Correctly (left), wrongly (middle) assigned, and unmatched (right) fractions as a function of the efficiency for events in the signal sample to have a jet which passes a $t \rightarrow sW$ score selection.

the fraction of signal events which pass a given selection on the primary s -jet score. Table III presents the fractions of events where the predicted primary s -jet matches each of the MC truth categories. The table shows that the baseline jet classifier is able to use the jet constituent information to reject more of the b -jets from $t \rightarrow bW$ than the high-level DiSaJa-H, while DiSaJa-L utilizes the whole event information to outperform the baseline in rejected $t \rightarrow bW$ b -jets. On the other hand, DiSaJa-H is better than the baseline model at

TABLE III. The fraction of $t\bar{t} \rightarrow sWbW$ events where the highest $t \rightarrow sW$ jet score is assigned to a given MC truth matched category. The possible jet-type categories we consider are the signal s -jet from the $t \rightarrow sW$ decay, the b -jet from the $t \rightarrow bW$ decay, or other, a jet in the event which does not match to a top decay.

	$t \rightarrow sW$	$t \rightarrow bW$	Other
Baseline	80.0%	2.8%	17.1%
DiSaJa-H	84.5%	4.8%	10.7%
DiSaJa-L	85.9%	1.8%	12.4%

rejecting jets not from a top decay, and though DiSaJa-L selects more nontop decay jets than DiSaJa-H, it is still the highest performing classifier.

Figure 7 shows the normalized distributions of the score for the signal and background processes using the baseline and DiSaJa models. Light jets are more predominant in the $DY + jj$ process than in the $t\bar{t}$ and single top processes, which are likely to contain mistagged b -jets. The score distribution for the light jet is quite flat across the discriminant variable; therefore, picking out the highest discriminant score from several light jets sculpts the distribution toward higher values of the discriminant. Therefore, the $DY + jj$ process behaves differently from other backgrounds, and is closer to the signal distribution. Figure 8 presents the receiver operating characteristic (ROC) curves and the area under the curve (AUC) values for the baseline and the DiSaJa models. It shows the $t\bar{t}$, DY , and single top background rejection rate corresponding to a given signal efficiency. The results demonstrate that the DiSaJa models achieve better assignments to $t \rightarrow sW$ than the baseline model. Additionally, DiSaJa-L outperforms DiSaJa-H.




FIG. 7. Normalized distributions of the highest $t \rightarrow sW$ category scores for jets in events across different classification methods. The left shows the distribution using the baseline model. The middle illustrates the distribution using DISAJA-H, while the right panel displays the distribution using DISAJA-L. All distributions are normalized to 1 for comparative purposes.




FIG. 8. ROC curves comparing the baseline and DISAJA models. The background rejection rates are shown for three background processes: $t\bar{t}$ (left), DY (middle), and single top (right). The AUC value for each model is given in the legends.

We perform statistical tests to evaluate model performance. For the tests, we use a binned profile likelihood fit using the CMS Combine framework [60]. The observable for the fit is the highest $t \rightarrow sW$ assignment score as shown in Fig. 9 and the parameter of interest (POI) is the signal strength μ scaling the signal yield defined as $\mu = \frac{|V_{ts}|^2}{|V_{ts}^{\text{PDG}}|^2}$, where $|V_{ts}^{\text{PDG}}| = 4.110 \times 10^{-2}$. Systematic uncertainties expected to have a sizable impact on this study are included when deriving the expected performance of the measurements, as well as the statistical uncertainty due to the simulation sample size. We include the b -tagging uncertainty by varying the relative b -tagging efficiency by 2.5% and 10% for b - and non- b -jets, respectively, which are the largest relative changes across the full p_T range shown in the CMS b -tagging study [43], thus giving a conservative uncertainty estimate. The jet energy scale uncertainty is taken into account by increasing or decreasing the p_T of each jet by 4%, which is taken from the jet energy scale and resolution measurement study by the CMS Collaboration [61–63].

Modeling uncertainties for the $t\bar{t}$ processes are also considered in the results. The uncertainties of the QCD scales at the matrix-element level are included by fixing the renormalization scale μ_R (factorization scale μ_F) while varying μ_F (μ_R) by a factor of 2 and 0.5 for up and down variations, respectively. Parton shower uncertainties are considered by varying the value of μ_R applied to the initial state radiation (ISR) or final state radiation (FSR) branchings by 2 and 0.5 for up and down variation, respectively. The PDF uncertainty is taken into account by following the PDF4LHC recommendation for Run 2 of the LHC Run 2 [64].

We calculate the expected significance of excluding $|V_{ts}| = 0$ using the test statistic $q = -2 \ln \frac{\mathcal{L}(0, \hat{\theta}_0)}{\mathcal{L}(\hat{\mu}, \hat{\theta})}$ where μ is the signal strength, θ are the nuisance parameters, and $\hat{\mu}$ and $\hat{\theta}$ are the unconstrained maximum likelihood estimator for μ and θ , respectively, and $\hat{\theta}_\mu$ is the value of the nuisance parameters that maximize the likelihood at a given μ . The test statistic is truncated to 0 when $\hat{\mu} < 0$. The expected significance is derived using the Asimov dataset and the




FIG. 9. Normalized score distribution for different models at an integrated luminosity of 138 fb^{-1} . The baseline model, DiSaJa-H, and DiSaJa-L are compared for signal-background separation.

asymptotic approximation of the profile likelihood ratio [65]. The calculation is performed on the observable distribution normalized to an integrated luminosity of 138 fb^{-1} , which is equivalent to the data collected at CMS during the LHC Run 2 period. Then, the expected significance is extrapolated up to 3000 fb^{-1} , which is expected to be collected during the upcoming HL-LHC experiment. Figure 10 illustrates a comparison of the significance for each model. The DiSaJa models outperform the baseline model, and the DiSaJa-L model performs better than the DiSaJa-H model. The DiSaJa-L model shows an expected exclusion significance greater than 3σ with the HL-LHC luminosity.

We calculate the expected CL_s upper limits at the 95% CL using the best-performing model, DiSaJa-L. For

the limit calculation, the test statistic $q = -2 \ln \frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\theta})}$ is employed. Depending on the value of $\hat{\mu}$, the test statistic is modified to $q = -2 \ln \frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(0, \hat{\theta}_0)}$ for $\hat{\mu} < 0$ and is set to $q = 0$ for $\hat{\mu} > \mu$ [60]. Using the Asimov dataset with $\mu = 0$, the expected median upper limit is derived and the expected $\pm 1\sigma$ and $\pm 2\sigma$ statistical fluctuations are extracted using asymptotic properties of the likelihood function [65], yielding $\mu < 0.7598^{+0.3058}_{-0.2171}$ based on the integrated luminosity of Run 2. The upper limit result is also projected to the Run 3 and HL-LHC luminosities and yields $\mu < 0.6699^{+0.2697}_{-0.1913}$ and $\mu < 0.5879^{+0.2367}_{-0.1679}$, respectively, as shown in Fig. 11.




FIG. 10. The expected significance of excluding scenarios with $|V_{ts}| = 0$ calculated for integrated luminosities from 138 fb^{-1} (CMS Run 2 luminosity) to 3000 fb^{-1} (HL-LHC luminosity). The significance is calculated by projecting the luminosity and takes into account systematic uncertainties.




FIG. 11. The expected CL_s upper limit on signal strength μ (DiSaJa-L). The expected CL_s upper limit is calculated from 138 fb^{-1} (CMS Run 2 luminosity) to 3000 fb^{-1} (HL-LHC luminosity) with luminosity projection and considering systematic uncertainties.




FIG. 12. Expected negative log-likelihood scan for the signal strength (μ) on an LHC Run 2 luminosity (138 fb^{-1}). The DiSaJa-L shows the best performance, providing bounds for $|V_{ts}|$ estimation.

We scan the negative log-likelihood ratio of each model, using the Asimov dataset with $\mu = 1$, as a function of the signal strength with the integrated luminosity of Run 2 to obtain the expected confidence interval for the measurement, as shown in Fig. 12. Among the models, DiSaJa-L shows the smallest interval, consistent with the other model comparisons above. With this model, the expected interval is $1.90 \times 10^{-2} < |V_{ts}| < 5.49 \times 10^{-2}$ at the 95% CL.

Lastly, we evaluate how strongly each uncertainty source affects the POI μ . For every source, the impact is defined as the shift of μ ($\Delta\mu$) obtained when the corresponding nuisance parameter is kept fixed at its postfit $\pm 1\sigma$, while all other nuisance parameters are profiled in the fit [60]. Table IV lists the uncertainties in descending order of their impact. The dominant contribution stems from the limited simulation sample size (as indicated by the MC statistics in the table), followed by b -tagging, QCD scale, FSR, JES,

TABLE IV. Impact of each uncertainty source on the signal strength μ . Per-bin MC statistical errors are reduced to MC statistics by the quadratic sum of them. The systematic effect is extracted based on the expected signal $\mu = 1$ and DiSaJa-L.

Source	Impact on μ ($\Delta\mu$)
MC statistics	$-0.341 / +0.344$
b -tagging	$-0.233 / +0.191$
QCD scale	$-0.190 / +0.185$
FSR	$-0.156 / +0.137$
JES	$-0.125 / +0.101$
ISR	$-0.0461 / +0.0695$
PDF	$-0.0216 / +0.0219$
Total	$-0.393 / +0.392$

ISR, and PDF effects. As the results are limited by the MC sample size statistical uncertainty, which is dominated by the $t\bar{t} \rightarrow bWbW$, the results we have shown are conservative estimates of the expected signal significance and limits. Our results show already that within the limits of the statistical uncertainty, the HL-LHC should be able to observe nonzero V_{ts} , and given collaboration-sized simulation resources, there is potential to increase the signal significance to the discovery level.

VI. CONCLUSION

We have studied the application of deep learning for the direct determination of $|V_{ts}|$ in the dileptonic $t\bar{t}$ final state events. The simulated sample reflects the environment of the CMS-like detector at the LHC experiment in the Run 2 period. Taking the primary s -jet tagging approach, we have developed a deep-learning-based jet discriminator, which we call DiSaJa-H, which uses as inputs all the high-level reconstructed objects of an event. The performance improvement using the DiSaJa-H method is tested by comparing it to the baseline model, and further performance improvement is achieved by using a model, DiSaJa-L, which uses the jet constituents as input variables. With the DiSaJa-L model, the expected significance for excluding $|V_{ts}| = 0$ assuming $|V_{ts}| = |V_{ts}^{\text{PDG}}|$ is 2.56σ , the median expected upper limit on μ at 95% CL is found to be 0.7598, and the confidence interval is $1.90 \times 10^{-2} < |V_{ts}| < 5.49 \times 10^{-2}$ at the 95% CL. Assuming the same collider environment as the Run 2 used in this paper, the statistical tests are extrapolated by projecting the integrated luminosity to 300 and 3000 fb^{-1} , corresponding to Run 3 and the HL-LHC, respectively. With the Run 3 projection, the results of the Run 2 are improved to 2.90σ for the expected significance and $\mu < 0.6699$ at the 95% CL is the median expected upper limit. With the HL-LHC projection, they are enhanced to 3.29σ and $\mu < 0.5879$. The DiSaJa models show a large performance increase over standard machine learning techniques, which will contribute significantly to measurement precision. Furthermore, the flexibility of the multidomain input and output of the model allows for it to be adapted to other analyses.

ACKNOWLEDGMENTS

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (Grants No. RS-2021-NR058866, No. RS-2023-NR076954, No. 2018R1C1B6005826, and No. RS-2021-NR058944), and the Basic Science Research Program through the NRF funded by the Ministry of Education (Grant No. 2018R1A6A1A06024977). This work was supported by the 2021 Research Fund from the University of Seoul.

J. H., W. J., and S. Y. contributed equally to this work.

DATA AVAILABILITY

The data that support the findings of this article are openly available [66].

-
- [1] M. Kobayashi and T. Maskawa, *CP* violation in the renormalizable theory of weak interaction, *Prog. Theor. Phys.* **49**, 652 (1973).
- [2] P. A. Zyla *et al.* (Particle Data Group), Review of particle physics, *Prog. Theor. Exp. Phys.* **2020**, 083C01 (2020).
- [3] J. Alwall, R. Frederix, J. M. Gerard, A. Giammanco, M. Herquet, S. Kalinin, E. Kou, V. Lemaître, and F. Maltoni, Is $V_{tb} \simeq 1$?, *Eur. Phys. J. C* **49**, 791 (2007).
- [4] A. Lenz and U. Nierste, Theoretical update of $B_s - \bar{B}_s$ mixing, *J. High Energy Phys.* **06** (2007) 072.
- [5] D. King, A. Lenz, and T. Rauh, B_s mixing observables and $|V_{td}/V_{ts}|$ from sum rules, *J. High Energy Phys.* **05** (2019) 034.
- [6] R. Aaij *et al.* (LHCb Collaboration), Precise determination of the $B_s^0 - \bar{B}_s^0$ oscillation frequency, *Nat. Phys.* **18**, 1 (2022).
- [7] Y. Aoki *et al.* (Flavour Lattice Averaging Group (FLAG) Collaboration), FLAG review 2021, *Eur. Phys. J. C* **82**, 869 (2022).
- [8] A. M. Sirunyan *et al.* (CMS Collaboration), Measurement of CKM matrix elements in single top quark t -channel production in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Phys. Lett. B* **808**, 135609 (2020).
- [9] G. Aad *et al.* (ATLAS Collaboration), Measurement of t -channel production of single top quarks and antiquarks in pp collisions at 13 TeV using the full ATLAS run 2 data sample, *J. High Energy Phys.* **05** (2024) 305.
- [10] A. Ali, F. Barreiro, and T. Lagouri, Prospects of measuring the CKM matrix element $|V_{ts}|$ at the LHC, *Phys. Lett. B* **693**, 44 (2010).
- [11] W. Jang, J. S. H. Lee, I. Park, and I. J. Watson, Measuring $|V_{ts}|$ directly using strange-quark tagging at the LHC, *J. Korean Phys. Soc.* **81**, 377 (2022).
- [12] D. A. Faroughy, J. F. Kamenik, M. Szwec, and J. Zupan, Accessing CKM suppressed top decays at the LHC, *SciPost Phys.* **16**, 131 (2024).
- [13] P. Abreu *et al.* (The DELPHI Collaboration), Measurement of the strange quark forward-backward asymmetry around the z^0 peak, *Eur. Phys. J. C* **14**, 613 (2000).
- [14] K. Abe *et al.* (The SLD Collaboration), First direct measurement of the parity-violating coupling of the z^0 to the s quark, *Phys. Rev. Lett.* **85**, 5059 (2000).
- [15] J. Duarte-Campderros, G. Perez, M. Schlaffer, and A. Soffer, Probing the Higgs–strange-quark coupling at e^+e^- colliders using light-jet flavor tagging, *Phys. Rev. D* **101**, 115005 (2020).
- [16] A. Albert *et al.*, Strange quark as a probe for new physics in the Higgs sector, in *Snowmass 2021* (2022), arXiv: 2203.07535.
- [17] J. Erdmann, A tagger for strange jets based on tracking information using long short-term memory, *J. Instrum.* **15**, P01021 (2020).
- [18] J. Erdmann, O. Nackenhorst, and S. V. Zeißner, Maximum performance of strange-jet tagging at hadron colliders, *J. Instrum.* **16**, P08039 (2021).
- [19] Y. Nakai, D. Shih, and S. Thomas, Strange jet tagging, arXiv:2003.09517.
- [20] F. Bedeschi, L. Gouskos, and M. Selvaggi, Jet flavour tagging for future colliders with fast simulation, *Eur. Phys. J. C* **82**, 646 (2022).
- [21] F. Blekman, F. Canelli, A. De Moor, K. Gautam, A. Ilg, A. Macchiolo, and E. Ploerer, Tagging more quark jet flavours at FCC-ee at 91 GeV with a transformer-based neural network, *Eur. Phys. J. C* **85**, 165 (2025).
- [22] Y. Kats and E. Ofir, From strange-quark tagging to fragmentation tagging with machine learning, *Phys. Rev. D* **111**, 034003 (2025).
- [23] J. S. H. Lee, I. Park, I. J. Watson, and S. Yang, Zero-permutation jet-parton assignment using a self-attention network, *J. Korean Phys. Soc.* **84**, 427 (2024).
- [24] J. Pearkes, W. Fedorko, A. Lister, and C. Gay, Jet constituents for deep neural network based top quark tagging, arXiv:1704.02124.
- [25] H. Qu, C. Li, and S. Qian, Particle transformer for jet tagging, in *Proceedings of the 39th International Conference on Machine Learning*, PMLR (2022). arXiv:2202.03772.
- [26] ATLAS Collaboration, Constituent-based top-quark tagging with the ATLAS detector, Technical Report, Report No. ATL-PHYS-PUB-2022-039, CERN, 2022.
- [27] ATLAS Collaboration, Constituent-based quark gluon tagging using transformers with the ATLAS detector, Technical Report, Report No. ATL-PHYS-PUB-2023-032, CERN, 2023.
- [28] J. A. Raine, M. Leigh, K. Zoch, and T. Golling, Fast and improved neutrino reconstruction in multineutrino final states with conditional normalizing flows, *Phys. Rev. D* **109**, 012005 (2024).
- [29] S. Qiu, S. Han, X. Ju, B. Nachman, and H. Wang, Holistic approach to predicting top quark kinematic properties with the covariant particle transformer, *Phys. Rev. D* **107**, 114029 (2023).
- [30] I. Zurbano Fernandez *et al.*, High-luminosity large hadron collider (HL-LHC): Technical design report, CERN Yellow Reports, [10.23731/CYRM-2020-0010](https://arxiv.org/abs/10.23731/CYRM-2020-0010) (2020), p. 10.
- [31] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [32] R. D. Ball *et al.* (NNPDF Collaboration), Parton distributions from high-precision collider data, *Eur. Phys. J. C* **77**, 663 (2017).

- [33] M. Czakon and A. Mitov, `Top++`: A program for the calculation of the top-pair cross-section at hadron colliders, *Comput. Phys. Commun.* **185**, 2930 (2014).
- [34] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to Pythia 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [35] T. Gehrmann, M. Grazzini, S. Kallweit, P. Maierhöfer, A. von Manteuffel, S. Pozzorini, D. Rathlev, and L. Tancredi, W^+W^- production at hadron colliders in next to next to leading order QCD, *Phys. Rev. Lett.* **113**, 212001 (2014).
- [36] M. Grazzini, S. Kallweit, D. Rathlev, and M. Wiesemann, $W^\pm Z$ production at hadron colliders in NNLO QCD, *Phys. Lett. B* **761**, 179 (2016).
- [37] F. Cascioli, T. Gehrmann, M. Grazzini, S. Kallweit, P. Maierhöfer, A. von Manteuffel, S. Pozzorini, D. Rathlev, L. Tancredi, and E. Weihs, ZZ production at hadron colliders in NNLO QCD, *Phys. Lett. B* **735**, 311 (2014).
- [38] R. Frederix and S. Frixione, Merging meets matching in MC@NLO, *J. High Energy Phys.* **12** (2012) 061.
- [39] A. M. Sirunyan *et al.* (CMS Collaboration), Extraction and validation of a new set of CMS Pythia8 tunes from underlying-event measurements, *Eur. Phys. J. C* **80**, 4 (2020).
- [40] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens, and M. Selvaggi (Delphes 3 Collaboration), Delphes 3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [41] M. Cacciari, G. P. Salam, and G. Soyez, Fastjet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [42] CMS Collaboration, Identification of b quark jets at the CMS experiment in the LHC run 2, Technical Report, Report No. CMS-PAS-BTV-15-001, CERN, Geneva, 2016.
- [43] A. M. Sirunyan *et al.* (CMS Collaboration), Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV, *J. Instrum.* **13**, P05011 (2017).
- [44] CMS Collaboration, Description and performance of track and primary-vertex reconstruction with the CMS tracker, *J. Instrum.* **9**, P10009 (2014).
- [45] A. M. Sirunyan *et al.* (CMS Collaboration), Measurements of $t\bar{t}$ differential cross sections in proton-proton collisions at $\sqrt{s} = 13$ TeV using events containing two leptons, *J. High Energy Phys.* **02** (2019) 149.
- [46] A. M. Sirunyan *et al.* (CMS Collaboration), Measurement of the $t\bar{t}$ production cross section, the top quark mass, and the strong coupling constant using dilepton events in pp collisions at $\sqrt{s} = 13$ TeV, *Eur. Phys. J. C* **79**, 368 (2019).
- [47] A. Tumasyan *et al.* (CMS Collaboration), Measurement of the top quark pole mass using $t\bar{t}$ + jet events in the dilepton final state in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **07** (2023) 077.
- [48] A. M. Sirunyan *et al.* (CMS Collaboration), Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC, *J. Instrum.* **16**, P05014 (2020).
- [49] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, Attention is all you need, in *Advances in Neural Information Processing Systems*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., Red Hook, New York, USA, 2017), Vol. 30.
- [50] D. Hendrycks and K. Gimpel, Gaussian error linear units (GELUs), [arXiv:1606.08415](https://arxiv.org/abs/1606.08415).
- [51] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [52] CMS Collaboration, Performance of quark/gluon discrimination in 8 TeV pp data, Technical Report, Report No. CMS-PAS-JME-13-002, CERN, Geneva, 2013.
- [53] T. Cornelis (CMS Collaboration), Quark-gluon jet discrimination at CMS, in 2nd Large Hadron Collider Physics Conference (2014), [arXiv:1409.3072](https://arxiv.org/abs/1409.3072).
- [54] R. D. Field and R. P. Feynman, A parametrization of the properties of quark jets, *Nucl. Phys.* **B136**, 1 (1978).
- [55] D. Krohn, M. D. Schwartz, T. Lin, and W. J. Waalewijn, Jet charge at the LHC, *Phys. Rev. Lett.* **110**, 212001 (2013).
- [56] W. J. Waalewijn, Calculating the charge of a jet, *Phys. Rev. D* **86**, 094030 (2012).
- [57] I. Loshchilov and F. Hutter, Decoupled weight decay regularization, [arXiv:1711.05101](https://arxiv.org/abs/1711.05101).
- [58] S. Watanabe, Tree-structured parzen estimator: Understanding its algorithm components and their roles for better empirical performance, [arXiv:2304.11127](https://arxiv.org/abs/2304.11127).
- [59] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, Optuna: A next-generation hyperparameter optimization framework, [arXiv:1907.10902](https://arxiv.org/abs/1907.10902).
- [60] A. Hayrapetyan *et al.* (CMS Collaboration), The CMS statistical analysis and combination tool: COMBINE, *Comput. Software Big Sci.* **8**, 19 (2024).
- [61] CMS Collaboration, Jet energy scale and resolution performance with 13 TeV data collected by CMS in 2016-2018, Technical Report, Report No. CMS-DP-2020-019, CERN, 2020.
- [62] CMS Collaboration, Jet energy scale and resolution measurement with run 2 legacy data collected by CMS at 13 TeV, Technical Report, Report No. CMS-DP-2021-033, CERN, 2021.
- [63] G. Agarwal (CMS Collaboration), Jet energy scale and resolution measurements in CMS, *Proc. Sci. ICHEP2022* (2022) 652 [[arXiv:2301.02175](https://arxiv.org/abs/2301.02175)].
- [64] J. Butterworth *et al.*, PDF4LHC recommendations for LHC run II, *J. Phys. G* **43**, 023001 (2016).
- [65] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, Asymptotic formulae for likelihood-based tests of new physics, *Eur. Phys. J. C* **71**, 1554 (2011); **73**, 2501(E) (2013).
- [66] J. Heo, W. Jang, J. S. H. Lee, Y. Jung Roh, I. James Watson, and S. Yang, Improving the direct determination of $|V_{ts}|$ using Deep Learning, [10.5281/zenodo.14862734](https://doi.org/10.5281/zenodo.14862734).