

Vector boson fusion signatures of superheavy Majorana neutrinos at muon colliders

Parham Dehghani^{1,*}, Mariana Frank^{1,†} and Benjamin Fuks^{2,‡}

¹*Department of Physics, Concordia University, 7141 Sherbrooke Street West, Montreal, Quebec H4B 1R6, Canada*

²*Laboratoire de Physique Théorique et Hautes Énergies (LPTHE), UMR 7589, Sorbonne Université et CNRS, 4 Place Jussieu, 75252 Paris Cedex 05, France*



(Received 13 June 2025; accepted 22 July 2025; published 13 August 2025)

We investigate the sensitivity of future high-energy muon colliders to heavy Majorana neutrinos, considering both opposite-sign ($\mu^+\mu^-$) and same-sign ($\mu^+\mu^+$) collision modes. We focus on $\mu^+\mu^-$ colliders operating at center-of-mass energies of 1, 3, and 10 TeV, as well as the proposed μ TRISTAN facility at 2 TeV, and we analyze the production of heavy neutrinos via vector boson fusion in the t -channel, a mechanism that becomes dominant in the multi-TeV regime. We evaluate its exclusion potential in terms of the heavy neutrino mass and the mixing of the heavy neutrino with its Standard Model counterparts, using both cut-based selections and boosted decision trees trained to exploit the distinct kinematic signatures of heavy Majorana neutrino exchanges. Our results demonstrate the complementarity between collider configurations and show that active-sterile mixing angles as small as 0.001 could be probed for neutrino masses up to 100 TeV, an experimentally inaccessible region of the parameter space at current facilities. Altogether, this work establishes the discovery potential of muon colliders for testing super-heavy Majorana neutrinos, complementary to conventional probes, and provides compelling motivation for the next generation of high-energy lepton colliders.

DOI: [10.1103/3sxx-glsw](https://doi.org/10.1103/3sxx-glsw)

I. INTRODUCTION

The discovery of neutrino oscillations stands as one of the most compelling pieces of evidence for physics beyond the Standard Model (BSM). Neutrinos were indeed long believed to be massless, in agreement with the original formulation of the Standard Model (SM) of particle physics. However, experimental anomalies hinted that neutrinos might behave in unexpected ways, which was conclusively confirmed by the Super-Kamiokande [1] and the Sudbury Neutrino Observatory experiments [2]. These provided clear evidence for neutrino oscillations from one flavor to another, which subsequently implied that neutrinos have nonzero masses, and that lepton flavor is not conserved during neutrino propagation in space-time.

These phenomena can be accommodated by introducing fields and interactions supplementing the SM, equipping

the theoretical description of nature with a neutrino mass model. The resulting class of BSM setups, collectively named seesaw models [3–16], include interactions of new charged or gauge-singlet fermions, scalar states with exotic quantum numbers and/or gauge bosons associated with new symmetries over a large range of masses and couplings. Searches for such signs of physics beyond the SM thus probe one of the most important questions in particle physics today and may open a window onto new sectors of the BSM theory.

In this work, we adopt a model-independent approach and explore BSM setups in which the neutrino sector is extended by one or more heavy neutral leptons, denoted generically as N , with masses in the GeV to multi-TeV range. These states are considered to mix with the SM active neutrinos through matrix elements $V_{\ell N}$, and could potentially be observed at present and future high-energy experiments. Here, we focus on the capabilities of proposed future muon-antimuon ($\mu^+\mu^-$) and muon-muon ($\mu^+\mu^+$) colliders designed to probe new physics at the multi-TeV scale, which are expected to offer unique sensitivity to detecting rare processes and heavy states thanks to their clean leptonic environment and high center-of-mass energy [17–25].

Several studies have shown that such facilities could significantly increase the sensitivity to heavy neutral

* Contact author: parham.dehghani@concordia.ca

† Contact author: mariana.frank@concordia.ca

‡ Contact author: fuks@lpthe.jussieu.fr

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.

leptons and lepton flavor violating processes. For example, in $\mu^+\mu^-$ collisions, a heavy neutrino N can be produced in association with a SM neutrino ν via t -channel W -boson or s -channel Z -boson exchanges. For processes where the heavy neutrino decays into a charged lepton and a W boson through its mixing with the SM, it has been demonstrated that with center-of-mass energies of 3 to 10 TeV and an integrated luminosity of 1 ab^{-1} searches can probe N masses between approximately 100 GeV and 10 TeV, for mixing angles $|V_{\ell N}|$ in the range of 10^{-2} to 10^{-3} [26–28]. By comparison, this reach extends significantly beyond that of current and proposed electron-positron and hadron colliders [29–45].

Complementary searches can be conducted at proposed μ^+e^+ colliders such as the μ TRISTAN facility. Although limited by a lower center-of-mass energy of 346 GeV, such setups could still probe mixing angles of an order a few times 10^{-3} for N masses above approximately 50 GeV and below the production threshold [46]. Further opportunities at μ TRISTAN arise from same-sign muon collisions ($\mu^+\mu^+$), which offer two promising production modes [47,48]. The first consists of direct N production, which is analogous to the $\mu^+\mu^-$ case and which gives access to heavy neutrino masses extending up to the collider’s center-of-mass energy, expected to be 1 or 2 TeV for mixings of approximately 10^{-3} . The second mode is an indirect channel where the heavy neutrino mediates a t -channel process leading to same-sign W^+W^+ pair production. This indirect signature allows sensitivity to N masses ranging up to approximately 1 TeV for small mixings around $|V_{\ell N}| \sim 0.02$ and up to about 100 TeV if the mixing increases to $|V_{\ell N}| \sim 0.2$. In addition, specific bounds have also been derived in complete neutrino mass models, leveraging the presence of extra states at the price of a not so clear interpretation due to a larger number of free parameters [49,50].

In this work, we explore an indirect yet powerful probe of heavy neutrinos, focusing both on same-sign and opposite-sign muon pair production via vector boson fusion (VBF) and mediated by t -channel exchanges of heavy neutrinos. By construction, this production mechanism, schematically described by the process $\mu^+\mu^\pm \rightarrow \nu_\mu \nu_\mu \mu^+\mu^\pm$ (with ν_μ standing for a generic neutrino and an antineutrino), has a unique sensitivity to new physics beyond the kinematic threshold of direct production. At the large hadron collider (LHC), same-sign dilepton signals via VBF have already been shown to provide constraints on heavy neutrinos with masses well above the collider’s hadronic center-of-mass energy [34,40,41,45,51]. This enhancement arises from the interplay between the high virtuality of the exchanged gauge bosons and the non-decoupling behavior of the heavy Majorana neutrino in the t -channel. It thus opens a complementary discovery path that is both less reliant on on-shell production and sensitive to the Majorana nature of the heavy states [52,53].

We simulate both the signal and the relevant SM backgrounds by means of state-of-the-art event generation, parton showering, and detector simulation tools in order to evaluate the projected sensitivity of several benchmark future collider configurations: a $\mu^+\mu^-$ collider operating at a center-of-mass energy of $\sqrt{s} = 10 \text{ TeV}$ with an integrated luminosity of 10 ab^{-1} , as well as lower-energy setups at $\sqrt{s} = 1$ and 3 TeV with 1 ab^{-1} . Moreover, we also consider the prospects for a $\mu^+\mu^+$ collider at $\sqrt{s} = 2 \text{ TeV}$ with 1 ab^{-1} . Our results demonstrate a remarkable signal and background separation power and that VBF-induced same-sign and opposite-sign muon production allows these machines to probe heavy neutrino masses far beyond the collider energy, reaching up to approximately 100 TeV for active-sterile mixing angles as small as $|V_{\mu N}| \sim 10^{-3}$. This significantly surpasses the reach of other production channels currently studied in the literature and highlights the exceptional potential of high-energy muon colliders as probes of seesaw-scale physics through indirect production mechanisms.

The remainder of this paper is organized as follows. In Sec. II, we introduce the simplified theoretical framework used to describe generic BSM extensions featuring a heavy neutrino (Sec. II A), along with the Monte Carlo simulation tool chain employed in our study (Sec. II B). We also detail our assumptions for the performance of a prospective muon collider detector and present a dedicated parametrization tailored to the simplified fast detector simulation (SFS) framework [54–56] of MADANALYSIS5 [57–59]. In Sec. III, we discuss our phenomenological analysis, considering both a cut-based selection (Sec. III A) and a multivariate one (Sec. III B). We finally summarize and conclude in Sec. IV.

II. THEORETICAL FRAMEWORK AND COMPUTATIONAL SETUP

A. A simplified model approach for heavy neutrino phenomenology

To assess the sensitivity of future colliders to heavy neutrinos through indirect VBF probes, we adopt a simplified model inspired by the type-I seesaw mechanism. In this framework, the SM is extended by a single sterile neutrino state ν_R which mixes with the SM active neutrino gauge eigenstates ν_{Le} , $\nu_{L\mu}$, and $\nu_{L\tau}$ [15,16]. The left-handed flavor-eigenstate neutrinos can then be written as

$$\nu_{Li} = \sum_{m=1}^3 U_{im} \nu_m + V_{iN} N, \quad (2.1)$$

where $i = e, \mu, \tau$. In this expression, the complex matrix elements U_{im} and V_{iN} parametrize the mixing between the interaction states ν_{Li} and the mass eigenstates of the light neutrinos ν_m and the heavy neutrino N , respectively. To

leading order in the active-sterile mixing elements $V_{\ell N}$, the heavy neutrino therefore inherits interactions with the electroweak sector. These include, in particular, charged-current couplings to the W boson,

$$\mathcal{L}_{\text{CC}} = -\frac{g}{\sqrt{2}} \sum_{\ell=e,\mu,\tau} \left(\sum_{i=1}^3 \bar{\nu}_i U_{\ell i}^* + \bar{N} V_{\ell N}^* \right) \not{W} P_L \ell + \text{H.c.}, \quad (2.2)$$

where g denotes the weak coupling constant, and P_L is the left-handed chirality projector. In the phenomenological approach followed in this work, we focus on the mixing between the heavy neutrino and the muon flavour, i.e., the mixing matrix element $V_{\mu N}$, which directly controls both the heavy neutrino production rate at muon colliders and its decay width. Rather than imposing constraints from a specific ultraviolet completion or a full seesaw model description, we treat $V_{\mu N}$ and the heavy neutrino mass M_N as free parameters, which allows us to explore the sensitivity of future muon colliders in a model-independent way while retaining the essential features of the seesaw-motivated dynamics.

Within this framework, we consider benchmark points in the parameter space ensuring a width-to-mass ratio Γ_N/M_N below 30% once all two-body decays into SM charged or neutral leptons and electroweak or Higgs bosons are included. This condition ensures that we remain in the regime where the simplified model is expected to yield reliable predictions. In addition, scenarios with $\Gamma_N/M_N \in [10, 30]\%$ are specifically highlighted, as cross section computations and event generation may suffer from reduced accuracy and require a more careful treatment of the heavy neutrino propagator.

In what follows, the heavy neutrino width is computed with MADWIDTH [60] and the HEAVYN [61] libraries built with FEYNRULES [62–64] and then exported to the UFO format [65,66]. The obtained predictions show that viable scenarios with neutrino masses above 1 TeV are constrained by the need to limit the value of the mixing matrix element $V_{\mu N}$ since the width scales proportionally to $|V_{\mu N}|^2$. This behavior is expected and reflects the enhanced decay rates induced by longitudinally polarized gauge bosons at high mass [15]. For instance, at $M_N = 1$ TeV, mixing values $|V_{\mu N}|^2 \lesssim 0.2$ keep the width-to-mass ratio below 30%, while at $M_N = 10$ TeV, the same condition requires $|V_{\mu N}|^2 \lesssim 0.002$.

B. Monte Carlo simulations

To investigate the considered VBF-induced heavy neutrino signal at future colliders, we employ a state-of-the-art Monte Carlo simulation chain. Hard-scattering signal and background events are generated at leading-order accuracy using the multipurpose event generator MADGRAPH5_AMC@NLO [67]. Whereas the SM implementation is used for the simulation of background events,

processes involving heavy neutrinos rely on the HEAVYN UFO libraries. The generated parton-level events are subsequently passed to PYTHIA8 [68] for parton showering (including QCD and QED radiation) and hadronization, while detector effects are accounted for using the SFS framework [54–56] embedded within the MADANALYSIS5 [57–59] platform that we also use for the implementation of our analysis.

For this purpose, we have developed a dedicated SFS detector card that approximates the performance characteristics expected of a future muon collider detector and relies on FASTJET [69,70] for event reconstruction. It incorporates smearing effects that arise predominantly from beam-induced backgrounds originating from in-flight muon decays, which are known to pose a significant challenge to object reconstruction [71–73]. The SFS card also integrates detector performance elements inspired by compact linear collider (CLIC) studies, such as robust tracking in the presence of beam-induced backgrounds [74,75], as well as design inputs from future circular collider infrastructure which is considered a viable option for muon production and acceleration [76,77].

We validate our SFS parametrization using inclusive top-antitop production in a muon collider environment, simulated via the process $\mu^+\mu^- \rightarrow t\bar{t}$ at a center-of-mass energy of $\sqrt{s} = 10$ TeV. This process is particularly suitable for validation, as it features a diverse set of final-state objects including electrons, muons, taus, photons, and jets after parton showering, hadronization, and reconstruction. To emulate realistic detector acceptance, we apply transverse momentum thresholds of 20 GeV for jets and taus and of 10 GeV for photons and leptons. We then assess the performance of our SFS setup by comparing the reconstructed object properties with those obtained from a more resource-intensive simulation chain based on DELPHES version 3.5 [78] and its muon collider card.¹ In this comparison, any reconstructed object with pseudorapidity $|\eta| > 2.5$ is discarded. Additionally, we impose isolation requirements: leptons must be separated from jets by $\Delta R_{\ell j} > 0.1$; otherwise, they are removed from the event and considered part of the jet. Isolated photons are similarly required to satisfy $\Delta R_{\gamma j} > 0.2$ with respect to all jets.

In our parametrization, the reconstruction efficiencies and resolution functions for electrons and muons are derived from the proposed performance of a CLIC-like detector [74]. For muons, we implement a constant reconstruction efficiency of 0.999 for energies above 2.5 GeV, and zero otherwise. For electrons, a more fine-grained efficiency function is used, depending on both the electron energy E and pseudorapidity η . For instance, the efficiency is taken to be 0.7 for central and low-energy electrons with $3 \text{ GeV} < E < 8 \text{ GeV}$ and $0.91 < |\eta| < 1.1$.

¹See the repository https://www.github.com/delphes/delphes/blob/master/cards/delphes_card_MuonColliderDet.tcl.




FIG. 1. Pseudorapidity (top row) and transverse momentum (bottom row) distributions of the leading electron (left) and muon (right). We compare Monte Carlo truth information (shaded orange) to results obtained using DELPHES (solid) and our SFS detector parametrization (dotted). The ratio to the truth distribution is shown in the lower panels of each figure.

The four-momenta of the reconstructed electrons and muons are then smeared using resolution functions that depend on the lepton transverse momentum p_T and pseudorapidity η and that vary with the lepton energy E . As an example, the resolution for a lepton with $2.0 < |\eta| < 2.1$ and $200 \text{ GeV} \leq E < 500 \text{ GeV}$ is set to a constant value of 6.276 MeV. For more details, we refer to the full SFS card implementation, available within MADANALYSIS5 from version 1.11.1.

The impact of the above performance choices is illustrated in Fig. 1, where we compare the reconstructed pseudorapidity (top row) and transverse momentum (bottom row) distributions of the leading electron (left) and muon (right). We show truth-level predictions (shaded orange), DELPHES-based results (solid), and

predictions obtained with our SFS implementation (dotted). The lower panels then display the ratio of the reconstructed distributions to the Monte Carlo truth, namely, the exact event information available after hadronization and before the simulation of any detector or reconstruction effect.

Overall, the two detector simulations yield similar lepton spectra, although some important differences emerge due to the distinct approaches to modeling detector effects. However, these differences are consistent with previous findings [54]. In particular, DELPHES-based results show a tendency for leptons to migrate toward the forward detector region, leading to an excess of reconstructed leptons compared to the truth-level prediction for pseudorapidities $1 \lesssim |\eta| \lesssim 2.5$. In contrast, this trend is less pronounced in

the SFS simulation, which more closely preserves the shape of the truth-level distribution.

Larger discrepancies are observed in the transverse momentum distributions. Both simulations produce similar high- p_T shapes, the overall efficiency being low due to reduced reconstruction of isolated leptons from boosted top decays. This effect is not surprising and naturally stems from the lepton isolation criteria imposed in the analysis. At low p_T , the SFS parametrization instead better reproduces the truth-level distribution. However, here events simulated using DELPHES tend to feature more reconstructed soft leptons that are simply discarded in the SFS case.

In our parametrization, jet reconstruction is performed using the anti- k_T algorithm [69] with a radius parameter $R = 0.4$. The impact of the detector includes, in particular, jet energy resolution effects, that we model through a smearing function inspired by the CLIC detector performance [74], which incorporates noise, stochastic, and calorimeter imperfection terms that depend on the jet energy and pseudorapidity. For example, in the central region ($|\eta| \leq 0.3$), the energy resolution is modeled as

$$\frac{\sigma}{E} = \sqrt{\left(\frac{1.38}{E}\right)^2 + \left(\frac{0.308}{\sqrt{E}}\right)^2} + 0.0025, \quad (2.3)$$

which yields an approximate resolution of 5.6% for a jet with $p_T = 100$ GeV. In addition, we apply an extra transverse momentum smearing of 2% for central jets ($|\eta| < 0.76$) and of 5% for more forward jets ($|\eta| \geq 0.76$).

To assess the impact of this jet modeling, we compare the reconstructed and truth-level distributions of the leading jet's transverse momentum (top) and pseudorapidity (bottom) in Fig. 2. As before, we contrast results from DELPHES (solid) and our SFS implementation (dotted) with the Monte Carlo truth (shaded orange). Both simulation approaches yield remarkably consistent predictions across the full p_T and $|\eta|$ ranges, in contrast with the lepton case discussed earlier where differences were more pronounced due to the treatment of object isolation. We emphasise that comprehensive details can be obtained from the SFS card available within MADANALYSIS5.

Additionally, for hadronic tau reconstruction we assume a constant tagging efficiency of 80% for taus with $p_T \geq 10$ GeV and 0% otherwise, with light jets being misidentified as hadronic taus with a probability of 2%. Furthermore, photon reconstruction performance is based on the electromagnetic calorimeter response used for electrons. In the central region of the detector ($|\eta| < 0.7$) the same parametrization is applied, while slight modifications are implemented in the forward region ($0.7 < |\eta| < 2.5$) as described in [74]. This results, for example, in a reconstruction efficiency of 94% for photons with $E \geq 2$ GeV in the central region and 90% in the forward region. The distributions of several properties associated with hadronic taus and photons show very good




FIG. 2. Transverse momentum (top) and pseudorapidity (bottom) distributions of the leading jet. We compare Monte Carlo truth information (shaded orange) with results from DELPHES (solid) and our SFS detector parametrization (dotted). The lower panels show the ratio to the truth-level distribution.

agreement between the DELPHES-based and SFS-based simulations and are therefore not shown here. As above, we refer to the SFS detector card included with MADANALYSIS5 version 1.11.1 for more detailed information on the detector modeling.

Before concluding this section, we present in Fig. 3 the distribution of the missing transverse energy (E_T), which is indirectly sensitive to the reconstruction performance of all visible objects in the event. We again compare Monte Carlo truth information (shaded orange) with the results obtained using DELPHES (solid black) and our SFS detector parametrization (dotted). The figure shows excellent overall agreement between the reconstructed-level and truth-level




FIG. 3. Missing transverse energy distribution derived from Monte Carlo truth information (shaded orange) and from predictions obtained using DELPHES (solid) and our SFS detector parametrization (dotted). The lower panels show the ratio of the reconstructed-level results to the truth-level distribution.

spectra across most of the E_T range, with the ratio to the Monte Carlo truth remaining close to unity. In the lower part of the spectrum ($E_T \in [200, 800]$ GeV), the SFS simulation predicts slightly more events than DELPHES, reflecting differences in the treatment of soft activity. At high E_T values (above 3 TeV), both detector simulations maintain equivalent reconstruction efficiency within the expected statistical uncertainties.

III. HEAVY NEUTRINOS AT FUTURE MUON COLLIDERS

To assess the sensitivity of future muon colliders to heavy Majorana neutrinos produced via vector boson fusion, we consider the processes $\mu^+\mu^- \rightarrow \mu^+\mu^-\nu_\mu\nu_\mu$ at a muon-antimuon collider and $\mu^+\mu^+ \rightarrow \mu^+\mu^+\nu_\mu\nu_\mu$ at an antimuon-antimuon collider such as the proposed μ TRISTAN machine, where we denote by ν_μ both a neutrino and an antineutrino. These 2-to-4 processes, that are potentially lepton-number violating, receive contributions from both s -channel and t -channel exchanges of heavy Majorana neutrinos, for which we provide representative diagrams in Fig. 4. These lead to a cross section scaling in $|V_{\mu N}|^4$ with the relative contributions of the s -channel and t -channel diagrams which depend strongly on the center-of-mass energy $\sqrt{s}$ and the heavy neutrino mass M_N . The former class of diagrams is dominant for N masses smaller than the center-of-mass energy, while the latter one only starts to dominate for much larger M_N values.

Our focus is particularly on the t -channel VBF topology, which, as shown in the LHC context [34,40,41,45,51], could provide sensitivity to neutrino masses well above the collider energy. Our study assumes integrated luminosities




FIG. 4. Representative Feynman diagrams for s -channel (top) and VBF t -channel (bottom) heavy neutrino contribution to the $\mu^+\mu^- \rightarrow \mu^+\mu^-\nu_\mu\nu_\mu$ process.

of 1, 1 and 10 ab^{-1} for $\mu^+\mu^-$ collisions at $\sqrt{s} = 1, 3$ and 10 TeV, respectively, and 1 ab^{-1} for $\mu^+\mu^+$ collisions at $\sqrt{s} = 2$ TeV. The dominant background arises from SM processes that lead to the same final-state topology as well as to a pair of muons, typically mediated by Z bosons, off-shell photons, and/or W bosons. For each collider setup, we design an analysis tailored to specific benchmark scenarios and then apply it over a broad range of heavy neutrino masses. Throughout it, we exploit the distinctive kinematic features of the signal to derive exclusion limits in the $(M_N, |V_{\mu N}|)$ parameter space, which we compare to existing bounds from the literature. In Sec. III A, we present a conventional cut-based analysis employing successive kinematic selections, while Sec. III B explores a multivariate approach based on boosted decision trees (BDTs), allowing for improved sensitivity in the targeted regions of parameter space.

A. A cut-based analysis of VBF-induced heavy neutrino production at muon colliders

We begin our study with the highest-energy muon collider currently under consideration, operating at a center-of-mass energy of $\sqrt{s} = 10$ TeV. To establish numerically robust sensitivity projections, we simulate one million signal events across a wide range of heavy neutrino masses spanning $M_N = 10$ GeV to 100 TeV, using the simulation framework described in Sec. II B, together with 10 millions background events. Mixing parameters




FIG. 5. Kinematic distributions used to define the selection cuts shown prior to the application of each respective cut. The panels show distributions for three benchmark signal scenarios corresponding to masses $M_N = 100, 500$, and 1000 GeV and mixing parameters $V_{\mu N} = 1.0, 1.0$, and 0.5 , alongside the dominant SM background contributions. We display the transverse momentum of the leading muon (top left), the missing transverse energy (top right), the azimuthal angle between the subleading muon and the missing transverse momentum (bottom left), and the sum of the pseudorapidities of the two leading muons (bottom right).

$V_{\mu N}$ are chosen such that the relative neutrino width Γ_N/M_N remains below 30% throughout.

All signal and background events are subject to baseline object identification criteria: leptons must be isolated from jets with a transverse separation $\Delta R_{\ell j} > 0.4$, and both jets and leptons are required to be central ($|\eta| < 2.5$) and energetic, with minimum transverse momenta of 20 GeV (jets) and 10 GeV (leptons). In particular, the centrality condition plays a crucial role in suppressing large-angle forward muons arising from SM processes such as $\mu^+\mu^- \rightarrow \mu^+\mu^-$. To optimize our event selection strategy, we perform a detailed kinematic analysis based on three benchmark scenarios with heavy neutrino masses $M_N = 100, 500$, and 1000 GeV and respective mixing parameters $V_{\mu N} = 1.0, 1.0$ and 0.5 . This set spans a broad region of the parameter space and allows us to design a selection strategy remaining efficient across the full range of interest. The final set of selection criteria is then defined to exploit the distinctive kinematic properties of VBF-induced heavy Majorana neutrino production.

The first selection step requires exactly two reconstructed opposite-sign muons, corresponding to the characteristic final-state topology of the signal. To suppress the SM di-muon background at high recoil energies, we impose a recoil mass constraint of $M_{\text{recoil}} < 9500$ GeV. Next, we exploit the asymmetric transverse momentum distribution of the signal: the leading muon is required to satisfy $p_T^{\mu_1} \in [1, 4]$ TeV while the subleading muon must have $p_T^{\mu_2} < 2.4$ TeV. We then require a minimum missing transverse energy of $E_T > 200$ GeV, consistent with the presence of energetic neutrinos in the final state. A tight azimuthal separation cut $\Delta\phi(\mu_2, \vec{p}_T) < 0.14$ further enhances the signal-to-background ratio by exploiting the strong angular correlation between the subleading muon and the missing momentum, a feature not typically present in SM backgrounds. The top left, top right, and bottom left panels of Fig. 5 illustrate the origin and effectiveness of these cuts, comparing signal and background distributions across all three benchmark points.

We further refine the event selection using pseudorapidity-based observables. The condition $|\eta_{\mu_1}| + |\eta_{\mu_2}| < 1.0$

TABLE I. Cut-flow table showing background and signal efficiencies for three benchmark scenarios (I, II, III) with heavy neutrino masses $M_N = 100, 500$, and 1000 GeV and respective mixing parameters $V_{\mu N} = 1.0, 1.0$, and 0.5 , at a $\sqrt{s} = 10$ TeV muon-antimuon collider. Initial event counts N_0 are normalized to an integrated luminosity of 10 ab^{-1} .

Selection	Background	Scenario I	Scenario II	Scenario III
	$N_0 = 5.32 \times 10^6$	$N_0 = 4.63 \times 10^9$	$N_0 = 4.63 \times 10^9$	$N_0 = 2.93 \times 10^8$
Initial	1	1	1	1
2 muons	0.901	0.995	0.995	0.994
$M_{\text{recoil}} < 9500 \text{ GeV}$	0.146	0.995	0.628	1.576
$p_T^{\mu_1} \in [1, 4] \text{ TeV}; p_T^{\mu_2} < 2.4 \text{ TeV}$	0.0180	0.745	0.744	0.729
$E_T > 200 \text{ GeV}$	0.0161	0.689	0.688	0.676
$\Delta\phi(\mu_2, \vec{p}_T) < 0.14$	0.00222	0.605	0.604	0.589
$ \eta_{\mu_1} + \eta_{\mu_2} < 1.0$	4.36×10^{-5}	0.335	0.333	0.320
$\eta_{\mu^+} - \eta_{\mu^-} < 0.6$	2.33×10^{-5}	0.273	0.271	0.261
$M_{\mu_1\mu_2} > 2500 \text{ GeV}$	7.70×10^{-6}	0.232	0.231	0.222

ensures that both muons are produced (very) centrally, another feature of the signal illustrated in the bottom right panel of Fig. 5. In addition, the difference in pseudorapidity $\eta_{\mu^+} - \eta_{\mu^-} < 0.6$ serves to suppress backgrounds from processes such as $\mu^+\mu^- \rightarrow \mu^+\mu^-\nu_\mu\nu_\mu$, and a final requirement on the di-muon invariant mass $M_{\mu_1\mu_2} > 2500 \text{ GeV}$ is imposed to isolate events consistent with the exchange of a heavy mediator.

The cut efficiencies for both signal and background are summarized in Table I. For an integrated luminosity of 10 ab^{-1} , the background is reduced to approximately 50 events, while retaining substantial signal yields, the exact value depending on $V_{\mu N}$ and m_N . The consistent performance of the cuts across different heavy neutrino masses validates the effectiveness of our analysis in capturing the general features of the VBF-induced heavy Majorana neutrino signal regardless of the heavy neutrino mass.


Building upon the analysis performed at $\sqrt{s} = 10 \text{ TeV}$, we extend our study to lower center-of-mass energies of 3 TeV and 1 TeV , assuming integrated luminosities of 1 ab^{-1} for both scenarios. Due to the modified kinematics and background cross sections at these reduced energies, we design dedicated sets of selection cuts analogous to those considered in the 10 TeV case, targeting the same benchmark scenarios. In particular, at lower center-of-mass energies, the heavy neutrino signal requires a relaxation of the pseudorapidity separation and missing transverse energy criteria compared to the 10 TeV setup. Likewise, the azimuthal correlation between the subleading muon and the missing transverse momentum necessitates a softer requirement. These revised kinematic selection criteria are summarized in the upper panel of Table II.

Finally, we investigate the prospects for heavy Majorana neutrino production via vector boson fusion at the $\mu\text{TRISTAN}$ collider, focusing on the same-sign $\mu^+\mu^+$ initial state. In our analysis, we consider an integrated

luminosity of 1 ab^{-1} and a center-of-mass energy of 2 TeV , consistent with the $\mu\text{TRISTAN}$ design parameters [22]. We target the process $\mu^+\mu^+ \rightarrow \mu^+\mu^+\nu_\mu\nu_\mu$, where we denote both neutrinos and antineutrinos generically by ν_μ , like for the opposite-sign collider case. Notably, the same-sign configuration only allows for the VBF t -channel topology, excluding s -channel diagrams that contribute in the opposite-sign scenario.

TABLE II. Selection criteria targeting the considered heavy neutrino signal at a future muon-antimuon collider expected to operate at $\sqrt{s} = 3 \text{ TeV}$ (top left) and 1 TeV (top right, and at a future same-sign $\mu^+\mu^+$ collider operating at $\sqrt{s} = 2 \text{ TeV}$ (bottom).

Selection	
$\sqrt{s} = 3 \text{ TeV}$	$\sqrt{s} = 1 \text{ TeV}$
2 muons	2 muons
$M_{\text{recoil}} < 2400 \text{ GeV}$	$M_{\text{recoil}} < 700 \text{ GeV}$
$p_T^{\mu_1} < 400 \text{ GeV}; p_T^{\mu_2} < 200 \text{ GeV}$	...
$E_T > 60 \text{ GeV}$	$E_T > 30 \text{ GeV}$
$\Delta\phi(\mu_2, \vec{p}_T) < 1.5$	$\Delta\phi(\mu_2, \vec{p}_T) > 0.14$
$\eta_{\mu^+} - \eta_{\mu^-} < 2.2$	$\eta_{\mu^+} - \eta_{\mu^-} < 1.0$
$M_{\mu_1\mu_2} > 240 \text{ GeV}$	$ \eta_{\mu_1} + \eta_{\mu_2} < 1.7$
Selection at $\sqrt{s} = 2 \text{ TeV}$	
2 same-sign muons	
$E_T > 40 \text{ GeV}$	
$M_{\text{recoil}} > 450 \text{ GeV}$	
$p_T^{\mu_1} > 350 \text{ GeV}$	
$\Delta\phi(\mu_2, \vec{p}_T) < 0.4$	
$ \eta_{\mu_1} + \eta_{\mu_2} < 1.7$	
$M_{\mu_1\mu_2} < 1500 \text{ GeV}$	

FIG. 6. Same as Fig. 5 but for the μ TRISTAN collider configuration.

Our analysis strategy mirrors the one developed for the considered muon-antimuon collider cases, with optimized cuts tailored to three benchmark scenarios featuring heavy neutrino masses of $M_N = 100, 500$ and 1000 GeV. The selection criteria are summarized in the lower panel of Table II, and representative kinematic distributions are shown in Fig. 6. Events are selected based on the presence of significant recoil energy carried by the invisible neutrinos, contrasting with the opposite-sign case where limiting this recoil helped to suppress the SM background. This reflects the dominant t -channel nature of the signal in the same-sign configuration, and for the same reason, a stringent cut is also applied on the transverse momentum of the leading muon, $p_T^{\mu_1} > 350$ GeV. Furthermore, the azimuthal correlation between the subleading muon and the missing transverse momentum displays a distinctive behavior, which we exploit to enhance signal sensitivity while significantly reducing background contamination. Additional discriminating power is finally achieved using the pseudorapidity and di-muon invariant mass distributions.

In order to evaluate the parameter space coverage offered by our various analyses, we select benchmark points for each collider configuration spanning the relevant heavy neutrino mass range while ensuring that the width-to-mass

ratio satisfies $\Gamma_N/M_N < 30\%$. For each benchmark, we compute the signal significance σ using the standard expression,

$$\sigma = \frac{S}{\sqrt{S + B + \sigma_B^2}}, \quad (3.1)$$

where S and B denote the number of signal and background events passing the entire selection, respectively, and σ_B^2 represents the background uncertainty which we vary from an idealized 0% to a conservative 20% of the background yield. Leveraging the quartic dependence of the signal cross section on $|V_{\mu N}|$, we then derive 95% confidence level (CL) exclusion limits in the $(M_N, |V_{\mu N}|^2)$ plane, as shown in Fig. 7. We present results for the four considered muon collider configurations: $\mu^+\mu^-$ collisions at $\sqrt{s} = 10$ TeV (red), 3 TeV (blue), and 1 TeV (purple) assuming integrated luminosities of 10 ab^{-1} , 1 ab^{-1} , and 1 ab^{-1} , respectively, as well as $\mu^+\mu^+$ collisions at a $\sqrt{s} = 2$ TeV μ TRISTAN (green) with a luminosity of 1 ab^{-1} . For each curve, the thickness of the exclusion contour reflects the variation in systematic background uncertainties.

Our results indicate that sensitivity to heavy neutrinos with masses up to 100 TeV is technically possible but only




FIG. 7. 95% CL exclusions expected from the four muon colliders considered, projected in the $(M_N, |V_{\mu N}|^2)$ plane. The red, blue, and purple contours correspond to 10, 1, and 1 ab^{-1} of $\mu^+\mu^-$ collisions at $\sqrt{s} = 10, 3$ and 1 TeV, respectively, while the green contours correspond 1 ab^{-1} of $\mu^+\mu^+$ collisions at $\sqrt{s} = 2$ TeV, with band thickness reflecting background systematic uncertainties varying from 0% to 20%. In addition, we display the most up-to-date limits from the LHC [41,79] (dashed and solid orange) and the μ TRISTAN projections of [48] for direct (dash-dotted pink) and indirect (dash-dotted teal) searches for heavy neutrinos, as well as the projections from [27] for conventional searches at $\mu^+\mu^-$ colliders (dotted green). Unphysical regions where $\Gamma_N/M_N > 0.3$ (dark gray) and intermediate regions where $0.1 < \Gamma_N/M_N < 0.3$ (light pink) are also shown.

in extreme cases featuring $\Gamma_N/M_N > 30\%$. More realistic scenarios with $\Gamma_N/M_N < 30\%$ remain accessible up to around 20 TeV, provided the active-sterile mixing is of order a few percent. Notably, all collider configurations demonstrate sensitivity in the very-heavy regime thanks to their differing background characteristics and the dedicated selection strategies that we have designed. Their impact becomes evident when comparing our limits not only with those from [48] (dash-dotted curves) and [27] (dotted contour), both of which relying on conventional search channels, but also with those of the LHC [41,79] (orange). We find that for each setup the VBF-induced production of heavy neutrinos provides a complementary probe to conventional direct and indirect searches, enabling sensitivity especially when $M_N \gg \sqrt{s}$. This complementarity is particularly evident for μ TRISTAN, where meaningful constraints are even achievable down to $M_N \sim 500$ GeV. Conversely, for masses below the center-of-mass energy, standard search strategies remain more effective, often outperforming the VBF-based limits by up to an order of magnitude depending on the heavy neutrino mass and collider configuration.

An important feature of our results is the change in the shape of the exclusion contours around $M_N \approx 0.5\text{--}1$ TeV. This transition marks the shift in the dominant production mode from s -channel diagrams to t -channel VBF topologies, which is particularly visible from the 3 TeV collider exclusion contour. Here, for M_N values below 500 GeV, s -channel production maintains favorable signal rates while for M_N values higher than 1 TeV, the signal turns out to be dominated by VBF production. In the transition region, interference between the two channels finally yields an

enhancement of the production cross section, implying a much more significant sensitivity. Such a behavior is also seen in the 1 TeV case, but in a more minimal way, while for the other setups there is no sizeable s -channel contribution. Beyond the 1 TeV mass threshold, VBF production ensures that the cross section remains sizeable, even up to $M_N \sim 30$ TeV. Interestingly, above about 10 TeV, lower-energy muon colliders can surpass the sensitivity of a 10 TeV muon collider, which underscores the superior efficiency of VBF mechanisms at high masses. It also highlights the nontrivial interplay between collider energy and production topology: a higher center-of-mass energy does not always guarantee better sensitivity. While the signal cross section decreases with increasing heavy neutrino mass, scaling as $|V_{\mu N}|^4/M_N^4$ for large M_N values [34], the interplay between the collider energy and the VBF topology exploited by our selection cuts leads to enhanced sensitivity in this regime. The strong background suppression achieved by the analysis ensures that even a small number of signal events can yield significant sensitivity. This effect is especially pronounced at lower collision energies, where the SM backgrounds are kinematically suppressed due to the absence of any dominant VBF contributions, unlike at $\sqrt{s} = 10$ TeV where the remaining background resembles the signal topology. As a result, for heavy neutrino masses above 10 TeV, lower-energy muon colliders could be expected to outperform a 10 TeV collider in terms of sensitivity. Finally, we note that our results remain robust against the inclusion of background systematic uncertainties up to 20%, with only minor degradation in sensitivity observed across the explored parameter space.

TABLE III. Optimized hyperparameters for BDT models trained on three heavy Majorana neutrino benchmark scenarios for 10 TeV opposite-sign and 2 TeV same-sign muon collider configurations. The scenarios I, II, III, IV, V, and VI are defined by $(M_N, V_{\mu N}) = (100 \text{ GeV}, 1.0), (5 \text{ TeV}, 0.1), (20 \text{ TeV}, 0.025), (500 \text{ GeV}, 0.9), (2 \text{ TeV}, 0.2),$ and $(10 \text{ TeV}, 0.05)$, respectively, and all models are trained using binary logistic regression [88] with the area under curve (AUC) as the evaluation metric.

Hyperparameter	Opposite sign			Same sign		
	I	II	III	IV	V	VI
max_depth	10	3	3	3	10	10
learning_rate	0.05	0.05	0.1	0.05	0.05	0.1
n_estimators	200	500	300	500	300	400
min_child_weight	5	1	3	1	3	2
gamma	0.1	0.1	0.2	0.1	0.2	0.2
subsample	0.80	0.95	0.95	0.95	0.80	0.80
colsample_bytree	0.8	0.8	0.9	0.8	0.8	0.8
alpha	0.4	0.2	0.4	0.2	0.2	0.2
early_stopping_rounds	20	20	20	20	20	20

B. Improving sensitivity to VBF-induced heavy neutrino production with BDTs

To refine our assessment of the reach of future muon colliders to heavy Majorana neutrinos, we rely on gradient boosting techniques to improve the 95% confidence level exclusions obtained from the traditional cut-based analyses implemented in Sec. III A. This approach allows us to exploit the full multidimensional space of observables used so far without the need for manual selections. Unlike simple sequential cut-based methods, boosted decision trees (BDTs) can hence capture nonlinear correlations between variables and better exploit the distinctive topological characteristics of a signal [80–86]. To quantify the potential improvements afforded by this technique in the model under study, we focus on two benchmark colliders: the most optimistic muon collider setup that we have considered, namely, $\mu^+\mu^-$ collisions at $\sqrt{s} = 10 \text{ TeV}$ with a luminosity of 10 ab^{-1} and the $\mu\text{TRISTAN}$ configuration corresponding to $\mu^+\mu^+$ collisions at $\sqrt{s} = 2 \text{ TeV}$ with a luminosity of 1 ab^{-1} . In both cases, we specifically target VBF-mediated heavy neutrino indirect production at very large masses.

For the $\mu^+\mu^-$ collider configuration, we consider three benchmark scenarios with $M_N = 100 \text{ GeV}$, 5 TeV , and 20 TeV and corresponding active-sterile mixing parameters set to $V_{\mu N} = 1.0, 0.1,$ and 0.025 , respectively. For the same-sign $\mu^+\mu^+$ case, we instead choose $M_N = 500 \text{ GeV}$, 2 TeV , and 10 TeV . For both signal and background samples, we apply the same baseline object selection criteria as defined in Sec. III A. In particular, leptons are required to be isolated from jets by $\Delta R_{\ell j} > 0.4$, and all jets and leptons must satisfy $|\eta| < 2.5$ with $p_T > 20 \text{ GeV}$ for jets and $p_T > 10 \text{ GeV}$ for leptons. To train the BDT classifier, we construct a comprehensive set of kinematic features for each event. These include the transverse momentum, pseudorapidity, and azimuthal angle of the

leading and subleading muons ($p_T^{\mu_1}, p_T^{\mu_2}, \eta_{\mu_1}, \eta_{\mu_2}, \phi_{\mu_1}, \phi_{\mu_2}$), as well as properties of the di-muon system such as its invariant mass $M_{\mu_1\mu_2}$, the azimuthal separation $\Delta\phi(\mu_1, \mu_2)$, and the recoil mass M_{recoil} . We further include derived quantities such as the transverse momentum ratio $p_T^{\mu_1}/p_T^{\mu_2}$, the sum $p_T^{\mu_1} + p_T^{\mu_2}$, the pseudorapidity sum $|\eta_{\mu_1}| + |\eta_{\mu_2}|$, the difference $\eta_{\mu^+} - \eta_{\mu^-}$, and the angular distance $\Delta R(\mu_1, \mu_2)$. In addition, we consider the missing transverse energy E_T , its angular separation from each muon $\Delta\phi(\mu_1, \vec{p}_T)$ and $\Delta\phi(\mu_2, \vec{p}_T)$, and the ratio $E_T/(p_T^{\mu_1} + p_T^{\mu_2})$. Although several of these variables are correlated, we show below that only a reduced subset of them significantly drives the sensitivity of the analysis.

Our BDT analysis is implemented using the XGBOOST package [87], which offers several advantages for high-energy physics applications. These include robust regularization, efficient handling of sparse data and parallel tree construction. The signal and background samples are split into training (70%) and test (30%) sets, using stratified sampling to preserve the signal-to-background ratio across subsets. To prevent overfitting, we further split the training data into a training subset (80%) and a validation subset (20%), enabling early stopping during training when the validation loss ceases to improve. Finally, the BDT hyperparameters are optimized to maximize the area under the receiver operating characteristic (ROC) curve for each of the considered signal scenarios.

To illustrate our findings, we summarize the resulting BDT configurations for each benchmark scenario in Table III, for the case of 10 TeV muon-antimuon collisions (second, third, and fourth columns) and $\mu^+\mu^+$ collisions at 2 TeV (last three columns). For the 10 TeV $\mu^+\mu^-$ configuration, the hyperparameter choices vary across benchmarks, reflecting different classification challenges




FIG. 8. Classification performance of BDT models trained for three representative heavy Majorana neutrino benchmarks with masses M_N fixed to 100 GeV, 5 TeV, and 20 TeV for $\mu^+\mu^-$ collisions at $\sqrt{s} = 10$ TeV (left) and 500 GeV, 2 TeV, and 10 TeV for $\mu^+\mu^+$ collisions at $\sqrt{s} = 2$ TeV (right). The performance is quantified using the area under the ROC curve, plotted as a function of the heavy neutrino mass M_N .

as the production mechanism shifts from mixed-channel to VBF-dominated regimes. For instance, the 100 GeV model uses significantly deeper trees (a maximum depth of 10), suggesting the need for more complex decision boundaries when both s -channel and t -channel contributions are relevant. Meanwhile, the number of estimators is largest in the intermediate 5 TeV mass scenario (500 estimators), which suggests that model complexity peaks in the intermediate mass regime. For the highest mass scenario, the learning rate and tree feature sampling increase to 0.1 and 0.9 respectively, which improves the training efficiency and stability of gradient descent in the presence of the large-momentum-transfer signatures typical of VBF-driven processes. The resulting BDT classifiers achieve area under curve (AUC) values of 0.999, 0.869, and 0.795 for the 100 GeV, 5 TeV, and 20 TeV benchmarks, respectively, demonstrating strong classification performance across a broad range of signal topologies. In contrast, for the 2 TeV $\mu^+\mu^+$ collider configuration, the hyperparameters are more similar across the three benchmark scenarios as the signal topology is more uniform. However, comparing the results in the last three columns of Table III highlights the more complex situation corresponding to the lightest scenario, yielding a signal harder to distinguish from the background. This is handled by requiring less deep trees with more estimators.

For simplicity, we select a single BDT model that can be used uniformly across the entire heavy neutrino mass range. We evaluate the classification performance of the three trained models on test samples spanning M_N values from 10 GeV to 100 TeV. Figure 8 shows the AUC as a function of M_N for $\mu^+\mu^-$ collisions at $\sqrt{s} = 10$ TeV (left) and $\mu^+\mu^+$ collisions at $\sqrt{s} = 2$ TeV (right). For the case of opposite-sign muon collisions, the best-performing model across most of the mass range is the one trained at $M_N = 5$ TeV. While the model trained at 100 GeV excels in the low-mass region ($M_N \lesssim 1$ TeV), its performance

drops significantly for heavier neutrino masses where VBF contributions dominate. Conversely, the 20 TeV-trained model captures the high-mass behavior reasonably well but performs poorly in the low-mass regime. The 5 TeV model subsequently offers a balanced performance across the entire mass range and successfully tracks the transition from s -channel to VBF topologies. It thus makes it a suitable compromise for a global analysis.

For same-sign muon collisions at $\sqrt{s} = 2$ TeV, the signal is VBF-dominated across the full mass range. Interestingly, in this case, the low-mass (500 GeV) model performs slightly better than the others, likely due to its training on VBF-like topologies already prominent at low M_N values requiring a different BDT model. Based on these observations, we adopt the 5 TeV BDT model for $\mu^+\mu^-$ collisions and the 500 GeV model for $\mu^+\mu^+$ collisions.

Next, we examine the relative contribution of each input feature to the BDT's classification power as quantified by the gain-based importance metric. As an example, Fig. 9 shows results for $\mu^+\mu^-$ collisions at $\sqrt{s} = 10$ TeV (left) and $\mu^+\mu^+$ collisions at $\sqrt{s} = 2$ TeV (right). Although many kinematic variables were provided to the model, only a few emerge as truly impactful. These features are found to be mostly uncorrelated, which allows the BDT to construct independent decision boundaries and avoid redundancy in the classification process. In the case of a 10 TeV opposite-sign muon collider, the azimuthal separation between the leading muons $\Delta\phi(\mu_1, \mu_2)$ stands out as the most discriminating variable, followed by the invariant mass $M_{\mu_1\mu_2}$ and the scalar sum of the muons transverse momenta $p_T^{\mu_1} + p_T^{\mu_2}$. Their ratio $p_T^{\mu_1}/p_T^{\mu_2}$, as well as their recoil M_{recoil} , contribute next, while still significantly. This ranking reflects the model's ability to capture the complementary angular and energy dependence of the signal, especially in the VBF-dominated regime. In contrast, pseudorapidity-based observables, critical in our cut-based




FIG. 9. Importance ranking calculated using the gain-based method implemented in XGBOOST, for the opposite-sign muon collider configuration at $\sqrt{s} = 10$ TeV (left) and the same-sign muon collider configuration at $\sqrt{s} = 2$ TeV (right).

approach to achieve acceptable background rejection, appear to contribute little to the multivariate classification. This can be attributed to correlations with other variables that make them redundant in the presence of more directly informative features. In contrast, for the 2 TeV same-sign muon collider, the same variables for the final-state muon pseudorapidities play a much more important role, the signal being here purely due to VBF-induced contributions, with the subleading variables being similar as in the opposite-sign case.

We remember that while including correlated variables may help improve raw sensitivity, they can complicate the estimation of systematic uncertainties as correlations are more difficult to propagate. In our case, however, only a

limited subset of largely uncorrelated input features significantly contributes to the discriminative power of the model. By efficiently exploiting orthogonal information across the input space, while down-weighting redundant variables, our BDT models can then be considered reliable for signal and background classification.

We now apply the two previously trained BDT models to derive projected exclusion limits in the $(M_N, |V_{\mu N}|^2)$ plane. After optimizing the selection thresholds for each point in this plane to maximise sensitivity, the exclusion contours are extracted following the procedure outlined in Sec. III A. The results are presented in Fig. 10, where we compare the reach of $\mu^+\mu^-$ collisions at $\sqrt{s} = 10$ TeV with a luminosity of 10 ab^{-1} (blue) to that of $\mu^+\mu^+$ collisions at $\sqrt{s} = 2$ TeV




FIG. 10. Projected 95% CL exclusion limits in the $(M_N, |V_{\mu N}|^2)$ plane using our two BDT-based analyses. The blue contour corresponds to $\mu^+\mu^-$ collisions at $\sqrt{s} = 10$ TeV with an integrated luminosity of 10 ab^{-1} , while the olive contour shows results for $\mu^+\mu^+$ collisions at $\sqrt{s} = 2$ TeV with 1 ab^{-1} . The thickness of the band reflects the impact of background systematic uncertainties, ranging from 0% to 20%. For comparison, existing LHC limits [41,79] (dashed and solid orange) are shown, along with direct (dash-dotted pink) and indirect (dash-dotted teal) projections from μ TRISTAN [48] and estimates from conventional $\mu^+\mu^-$ searches at future colliders [27] (dotted green). Unphysical regions where the heavy neutrino width exceeds 30% of its mass ($\Gamma_N/M_N > 0.3$) are shaded in dark gray, while intermediate-width regions ($0.1 < \Gamma_N/M_N < 0.3$) are shown in light pink.

with 1 ab^{-1} (olive). Compared to the cut-based approach (see Fig. 7), the BDT analysis leads to a substantial sensitivity gain, particularly in the heavy mass regime with $M_N \gtrsim 10 \text{ TeV}$ and for a collider operating at $\sqrt{s} = 10 \text{ TeV}$. In this region, the signal is dominated by VBF-induced production, which the BDT is especially well suited to identify thanks to its ability to exploit multidimensional correlations in the final-state kinematics. In particular, Fig. 9 shows that the nonlinear correlations among the $\Delta\phi(\mu_1, \mu_2)$, $M_{\mu_1\mu_2}$ and $p_T^{\mu_1} + p_T^{\mu_2}$ variables play a key role in enhancing the sensitivity; those correlations being otherwise inaccessible through sequential cuts. Moreover, the fixed selection criteria used in the cut-based analysis were optimized for specific benchmark scenarios and thus become increasingly suboptimal at higher M_N values where the signal kinematics deviate significantly from those benchmarks. The BDT analysis hence extends the reach down to $|V_{\mu N}|^2 \sim 10^{-6}$ for the largest masses considered, improving the results over those of cut-based limits by about two orders of magnitude.

IV. SUMMARY AND CONCLUSIONS

In this study, we explored the sensitivity of future muon colliders to heavy Majorana neutrinos, focusing on both opposite-sign ($\mu^+\mu^-$) and same-sign ($\mu^+\mu^+$) configurations. We considered realistic collider scenarios, including center-of-mass energies of 1, 3, and 10 TeV for $\mu^+\mu^-$ machines and 2 TeV for the proposed $\mu\text{TRISTAN}$ facility. Our analysis then leveraged VBF processes in the t -channel as the dominant production mechanism of Majorana neutrinos. Using both cut-based and multivariate (BDT-based) analysis techniques, we assessed the exclusion potential in the $(M_N, |V_{\mu N}|^2)$ plane on the basis of state-of-the-art simulations, made possible through the development and validation of an SFS muon collider detector card for the MADANALYSIS5 framework.

We found that future muon colliders could probe heavy neutrinos with masses up to 100 TeV, well beyond the reach

of current experiments, and active-sterile mixings as small as $|V_{\mu N}|^2 \sim 10^{-6}$ in the most favorable benchmark scenarios. Furthermore, our results highlight a strong complementarity between collider types: while same-sign $\mu^+\mu^+$ collisions isolate the t -channel contribution cleanly, opposite-sign $\mu^+\mu^-$ setups benefit from both s -channel and t -channel interference and a broader kinematic phase space. In particular, our BDT analysis proves powerful at high energies, fully exploiting the angular and energy correlations induced by heavy Majorana t -channel exchange.

Altogether, our findings underscore the unique discovery potential of high-energy opposite-sign and same-sign muon colliders to probe the Majorana nature of neutrinos across a vast mass range, extending the scope of earlier studies available from the literature by targeting the super-heavy mass regime. This therefore provides strong theoretical motivation for their continued development, with the prospect of shedding light on one of the most pressing open questions in particle physics, the origins of neutrino masses.

ACKNOWLEDGMENTS

The authors are grateful to Arindam Das for enlightening comments on the manuscript. The work of M. F. has been partly supported by NSERC through Grant No. SAP105354 and that of B. F. by Grant No. ANR-21-CE31-0013 (project DMwithLLPatLHC) from the French *Agence Nationale de la Recherche*.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

-
- [1] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), Evidence for oscillation of atmospheric neutrinos, *Phys. Rev. Lett.* **81**, 1562 (1998).
 - [2] Q. R. Ahmad *et al.* (SNO Collaboration), Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory, *Phys. Rev. Lett.* **89**, 011301 (2002).
 - [3] P. Minkowski, $\mu \rightarrow e\gamma$ at a rate of one out of 10^9 muon decays?, *Phys. Lett.* **67B**, 421 (1977).
 - [4] M. Gell-Mann, P. Ramond, and R. Slansky, Complex spinors and unified theories, *Conf. Proc. C* **790927**, 315 (1979).
 - [5] S. L. Glashow, The future of elementary particle physics, *NATO Sci. Ser. B* **61**, 687 (1980).
 - [6] R. N. Mohapatra and G. Senjanovic, Neutrino mass and spontaneous parity nonconservation, *Phys. Rev. Lett.* **44**, 912 (1980).
 - [7] T. Yanagida, Horizontal symmetry and masses of neutrinos, *Prog. Theor. Phys.* **64**, 1103 (1980).
 - [8] M. Magg and C. Wetterich, Neutrino mass problem and gauge hierarchy, *Phys. Lett.* **94B**, 61 (1980).
 - [9] R. E. Shrock, General theory of weak leptonic and semi-leptonic decays. 1. Leptonic pseudoscalar meson decays,

- with associated tests for, and bounds on, neutrino masses and lepton mixing, *Phys. Rev. D* **24**, 1232 (1981).
- [10] J. Schechter and J. W. F. Valle, Neutrino masses in $SU(2) \times U(1)$ theories, *Phys. Rev. D* **22**, 2227 (1980).
- [11] T. P. Cheng and L.-F. Li, Neutrino masses, mixings and oscillations in $SU(2) \times U(1)$ models of electroweak interactions, *Phys. Rev. D* **22**, 2860 (1980).
- [12] R. N. Mohapatra and G. Senjanovic, Neutrino masses and mixings in gauge models with spontaneous parity violation, *Phys. Rev. D* **23**, 165 (1981).
- [13] G. Lazarides, Q. Shafi, and C. Wetterich, Proton lifetime and fermion masses in an $SO(10)$ model, *Nucl. Phys. B* **181**, 287 (1981).
- [14] R. Foot, H. Lew, X. G. He, and G. C. Joshi, Seesaw neutrino masses induced by a triplet of leptons, *Z. Phys. C* **44**, 441 (1989).
- [15] A. Atre, T. Han, S. Pascoli, and B. Zhang, The search for heavy Majorana neutrinos, *J. High Energy Phys.* **05** (2009) 030.
- [16] Y. Cai, T. Han, T. Li, and R. Ruiz, Lepton number violation: Seesaw models and their collider tests, *Front. Phys.* **6**, 40 (2018).
- [17] M. Boscolo, J.-P. Delahaye, and M. Palmer, The future prospects of muon colliders and neutrino factories, *Rev. Accel. Sci. Technol.* **10**, 189 (2019).
- [18] J. P. Delahaye, M. Diemoz, K. Long, B. Mansoulié, N. Pastrone, L. Rivkin, D. Schulte, A. Skrinsky, and A. Wulzer, Muon colliders, *arXiv:1901.06150*.
- [19] K. Long, D. Lucchesi, M. Palmer, N. Pastrone, D. Schulte, and V. Shiltsev, Muon colliders to expand frontiers of particle physics, *Nat. Phys.* **17**, 289 (2021).
- [20] H. Al Ali *et al.*, The muon Smasher's guide, *Rep. Prog. Phys.* **85**, 084201 (2022).
- [21] K. M. Black *et al.*, Muon collider forum report, *J. Instrum.* **19**, T02015 (2024).
- [22] Y. Hamada, R. Kitano, R. Matsudo, H. Takaura, and M. Yoshida, μ TRISTAN, *Prog. Theor. Exp. Phys.* **2022**, 053B02 (2022).
- [23] C. Accettura *et al.*, Towards a muon collider, *Eur. Phys. J. C* **83**, 864 (2023); **84**, 36(E) (2024).
- [24] K. Fridell, R. Kitano, and R. Takai, Lepton flavor physics at $\mu^+\mu^+$ colliders, *J. High Energy Phys.* **06** (2023) 086.
- [25] D. Akturk, B. Dagli, B. Ketenoglu, A. Ozturk, and S. Sultansoy, μ TRISTAN and LHC/Tevatron/FCC/SppC based antimuon-hadron colliders, *arXiv:2406.02647*.
- [26] T. Li, C.-Y. Yao, and M. Yuan, Searching for heavy neutral lepton and lepton number violation through VBS at high-energy muon colliders, *J. High Energy Phys.* **09** (2023) 131.
- [27] P. Li, Z. Liu, and K.-F. Lyu, Heavy neutral leptons at muon colliders, *J. High Energy Phys.* **03** (2023) 231.
- [28] T. H. Kwok, L. Li, T. Liu, and A. Rock, Searching for heavy neutral leptons at a future muon collider, *Phys. Rev. D* **110**, 075009 (2024).
- [29] W.-Y. Keung and G. Senjanovic, Majorana neutrinos and the production of the right-handed charged gauge boson, *Phys. Rev. Lett.* **50**, 1427 (1983).
- [30] F. del Aguila, J. A. Aguilar-Saavedra, and R. Pittau, Heavy neutrino signals at large hadron colliders, *J. High Energy Phys.* **10** (2007) 047.
- [31] S. Banerjee, P. S. B. Dev, A. Ibarra, T. Mandal, and M. Mitra, Prospects of heavy neutrino searches at future lepton colliders, *Phys. Rev. D* **92**, 075002 (2015).
- [32] M. Nemevšek, F. Nesti, and G. Popara, Keung-Senjanović process at the LHC: From lepton number violation to displaced vertices to invisible decays, *Phys. Rev. D* **97**, 115018 (2018).
- [33] S. Pascoli, R. Ruiz, and C. Weiland, Heavy neutrinos with dynamic jet vetoes: Multilepton searches at $\sqrt{s} = 14, 27$, and 100 TeV, *J. High Energy Phys.* **06** (2019) 049.
- [34] B. Fuks, J. Neundorff, K. Peters, R. Ruiz, and M. Saimpert, Majorana neutrinos in same-sign $W^\pm W^\pm$ scattering at the LHC: Breaking the TeV barrier, *Phys. Rev. D* **103**, 055005 (2021).
- [35] K. Mękała, J. Reuter, and A. F. Żarnecki, Heavy neutrinos at future linear e^+e^- colliders, *J. High Energy Phys.* **06** (2022) 010.
- [36] M. Nemevšek and F. Nesti, Left-right symmetry at an FCC-hh, *Phys. Rev. D* **108**, 015030 (2023).
- [37] M. Frank, B. Fuks, A. Jueid, S. Moretti, and O. Ozdal, A novel search strategy for right-handed charged gauge bosons at the Large Hadron Collider, *J. High Energy Phys.* **02** (2024) 150.
- [38] M. Aaboud *et al.* (ATLAS Collaboration), Search for heavy Majorana or Dirac neutrinos and right-handed W gauge bosons in final states with two charged leptons and two jets at $\sqrt{s} = 13$ TeV with the ATLAS detector, *J. High Energy Phys.* **01** (2019) 016.
- [39] M. Aaboud *et al.* (ATLAS Collaboration), Search for a right-handed gauge boson decaying into a high-momentum heavy neutrino and a charged lepton in pp collisions with the ATLAS detector at $\sqrt{s} = 13$ TeV, *Phys. Lett. B* **798**, 134942 (2019).
- [40] G. Aad *et al.* (ATLAS Collaboration), Search for Majorana neutrinos in same-sign WW scattering events from pp collisions at $\sqrt{s} = 13$ TeV, *Eur. Phys. J. C* **83**, 824 (2023).
- [41] G. Aad *et al.* (ATLAS Collaboration), Search for heavy Majorana neutrinos in $e \pm e \pm$ and $e \pm \mu \pm$ final states via WW scattering in pp collisions at $s=13$ TeV with the ATLAS detector, *Phys. Lett. B* **856**, 138865 (2024).
- [42] A. M. Sirunyan *et al.* (CMS Collaboration), Search for a heavy right-handed W boson and a heavy neutrino in events with two same-flavor leptons and two jets at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **05** (2018) 148.
- [43] A. M. Sirunyan *et al.* (CMS Collaboration), Search for heavy Majorana neutrinos in same-sign dilepton channels in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **01** (2019) 122.
- [44] A. Tumasyan *et al.* (CMS Collaboration), Search for a right-handed W boson and a heavy neutrino in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **04** (2022) 047.
- [45] A. Tumasyan *et al.* (CMS Collaboration), Probing heavy Majorana neutrinos and the Weinberg operator through vector boson fusion processes in proton-proton collisions at $s = 13$ TeV, *Phys. Rev. Lett.* **131**, 011803 (2023).
- [46] A. Das, J. Li, S. Mandal, T. Nomura, and R. Zhang, Testing tree level TeV scale type-I and type-II seesaw scenarios in μ TRISTAN, *arXiv:2410.21956*.

- [47] R. Jiang, T. Yang, S. Qian, Y. Ban, J. Li, Z. You, and Q. Li, Searching for Majorana neutrinos at a same-sign muon collider, *Phys. Rev. D* **109**, 035020 (2024).
- [48] C. H. de Lima, D. McKeen, J. N. Ng, M. Shamma, and D. Tuckler, Probing lepton number violation at same-sign lepton colliders, *Phys. Rev. D* **111**, 075002 (2025).
- [49] P. S. B. Dev, J. Heeck, and A. Thapa, Neutrino mass models at μ TRISTAN, *Eur. Phys. J. C* **84**, 148 (2024).
- [50] P. Bandyopadhyay and S. Parashar, Probing a scalar singlet-triplet extension of the standard model via vector boson fusion at a muon collider, *Phys. Rev. D* **110**, 115032 (2024).
- [51] B. Fuks, J. Neundorff, K. Peters, R. Ruiz, and M. Saimpert, Probing the Weinberg operator at colliders, *Phys. Rev. D* **103**, 115014 (2021).
- [52] P. S. B. Dev, A. Pilaftsis, and U.-k. Yang, New production mechanism for heavy neutrinos at the LHC, *Phys. Rev. Lett.* **112**, 081801 (2014).
- [53] D. Alva, T. Han, and R. Ruiz, Heavy Majorana neutrinos from $W\gamma$ fusion at hadron colliders, *J. High Energy Phys.* **02** (2015) 072.
- [54] J. Y. Araz, B. Fuks, and G. Polykratis, Simplified fast detector simulation in MADANALYSIS5, *Eur. Phys. J. C* **81**, 329 (2021).
- [55] J. Y. Araz, B. Fuks, M. D. Goodsell, and M. Utsch, Recasting LHC searches for long-lived particles with MADANALYSIS5, *Eur. Phys. J. C* **82**, 597 (2022).
- [56] J. Y. Araz, A. Buckley, and B. Fuks, Searches for new physics with boosted top quarks in the MADANALYSIS5 and Rivet frameworks, *Eur. Phys. J. C* **83**, 664 (2023).
- [57] E. Conte, B. Fuks, and G. Serret, MADANALYSIS5, A user-friendly framework for collider phenomenology, *Comput. Phys. Commun.* **184**, 222 (2013).
- [58] E. Conte, B. Dumont, B. Fuks, and C. Wymant, Designing and recasting LHC analyses with MADANALYSIS5, *Eur. Phys. J. C* **74**, 3103 (2014).
- [59] E. Conte and B. Fuks, Confronting new physics theories to LHC data with MADANALYSIS5, *Int. J. Mod. Phys. A* **33**, 1830027 (2018).
- [60] J. Alwall, C. Duhr, B. Fuks, O. Mattelaer, D. G. Öztürk, and C.-H. Shen, Computing decay rates for new physics theories with FEYNRULES and MADGRAPH5_AMC@NLO, *Comput. Phys. Commun.* **197**, 312 (2015).
- [61] C. Degrande, O. Mattelaer, R. Ruiz, and J. Turner, Fully-automated precision predictions for heavy neutrino production mechanisms at hadron colliders, *Phys. Rev. D* **94**, 053002 (2016).
- [62] N. D. Christensen, P. de Aquino, C. Degrande, C. Duhr, B. Fuks, M. Herquet, F. Maltoni, and S. Schumann, A comprehensive approach to new physics simulations, *Eur. Phys. J. C* **71**, 1541 (2011).
- [63] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, FEYNRULES2.0—A complete toolbox for tree-level phenomenology, *Comput. Phys. Commun.* **185**, 2250 (2014).
- [64] C. Degrande, Automatic evaluation of UV and R2 terms for beyond the standard model Lagrangians: A proof-of-principle, *Comput. Phys. Commun.* **197**, 239 (2015).
- [65] C. Degrande, C. Duhr, B. Fuks, D. Grellscheid, O. Mattelaer, and T. Reiter, UFO—The Universal FEYNRULES Output, *Comput. Phys. Commun.* **183**, 1201 (2012).
- [66] L. Darmé *et al.*, ‘UFO 2.0: The Universal Feynman Output’ format, *Eur. Phys. J. C* **83**, 631 (2023).
- [67] 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.
- [68] C. Bierlich *et al.*, A comprehensive guide to the physics and usage of PYTHIA8.3, *SciPost Phys. Codebases* **2022**, 8 (2022).
- [69] M. Cacciari, G. P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [70] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [71] F. Collamati, C. Curatolo, D. Lucchesi, A. Mereghetti, N. Mokhov, M. Palmer, and P. Sala, Advanced assessment of beam-induced background at a muon collider, *J. Instrum.* **16**, P11009 (2021).
- [72] C. Aimè, C. Riccardi, P. Salvini, and I. Vai (Muon Collider Collaboration), Simulated performance of a multi-purpose experiment at a muon collider, *Proc. Sci., ICHEP2022* (2022) 1109.
- [73] N. Bartosik *et al.* (Muon Collider Collaboration), Simulated detector performance at the muon collider, *arXiv:2203.07964*.
- [74] D. Arominski *et al.* (CLICdp Collaboration), A detector for CLIC: Main parameters and performance, *arXiv:1812.07337*.
- [75] C. Greife, Detector optimization studies and light Higgs decay into muons at CLIC, Ph.D. thesis, Bonn U., 2012, *arXiv:1402.2780*.
- [76] F. Zimmermann, LHC/FCC-based muon colliders, *J. Phys. Conf. Ser.* **1067**, 022017 (2018).
- [77] M. Benedikt *et al.*, Future circular hadron collider FCC-hh: Overview and status, *arXiv:2203.07804*.
- [78] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi (DELPHES3 Collaboration), DELPHES3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [79] A. Hayrapetyan *et al.* (CMS Collaboration), Search for heavy neutral leptons in final states with electrons, muons, and hadronically decaying tau leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **06** (2024) 123.
- [80] B. P. Roe, H.-J. Yang, J. Zhu, Y. Liu, I. Stancu, and G. McGregor, Boosted decision trees, an alternative to artificial neural networks, *Nucl. Instrum. Methods Phys. Res., Sect. A* **543**, 577 (2005).
- [81] Y. Coadou, Boosted decision trees and applications, *EPJ Web Conf.* **55**, 02004 (2013).
- [82] K. Albertsson *et al.*, Machine learning in high energy physics community white paper, *J. Phys. Conf. Ser.* **1085**, 022008 (2018).

- [83] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [84] A. S. Cornell, W. Doorsamy, B. Fuks, G. Harmsen, and L. Mason, Boosted decision trees in the era of new physics: A smuon analysis case study, *J. High Energy Phys.* **04** (2022) 015.
- [85] A. Choudhury, A. Mondal, and S. Sarkar, Searches for the BSM scenarios at the LHC using decision tree-based machine learning algorithms: A comparative study and review of random forest, AdaBoost, XGBoost and LightGBM frameworks, *Eur. Phys. J. Special Topics* **233**, 2425 (2024).
- [86] A. S. Cornell, B. Fuks, M. D. Goodsell, and A. M. Ncube, Improving smuon searches with neural networks, *Eur. Phys. J. C* **85**, 51 (2025).
- [87] T. Chen and C. Guestrin, XGBoost: A Scalable Tree Boosting System, in *KDD '16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016), 10.1145/2939672.2939785.
- [88] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning* (Springer, New York, 2009), 10.1007/978-0-387-84858-7.