

Data-driven predictions for dark photon and millicharged particle production

Elizabeth Allison^{1,2} and Nikita Blinov^{2,*}

¹*Neuroscience Graduate Program, Schulich School of Medicine and Dentistry,
Western University, London, Ontario, Canada*

²*Department of Physics and Astronomy, York University,
Toronto, Ontario, Canada*



(Received 12 December 2025; accepted 1 May 2026; published 22 May 2026)

Accurate signal predictions are essential for interpreting and optimizing fixed-target searches for new physics. Even in minimal models such as the dark photon (A') or millicharged particles (mCPs), theoretical uncertainties in hadronic production can be substantial. We introduce a data-driven framework that predicts both the rate and kinematic distributions of A' and mCP production directly from measured dilepton events, without relying on specific theoretical production models. This method uses the close correspondence between amplitudes for emission of A' or mCPs, and for off-shell Standard Model photon production, the latter being experimentally measurable in full differential form. We demonstrate that normalizing flow models can learn these distributions from data and serve as a fast, realistic Monte Carlo generator for dark sector signal simulations.

DOI: [10.1103/vzmz-jbck](https://doi.org/10.1103/vzmz-jbck)

I. INTRODUCTION

Fixed-target experiments are a powerful way of searching for light, weakly coupled particles in a variety of dark sector models [1,2]. Interpreting these searches, however, requires accurate predictions for both the production rate and kinematic distributions of potential signals. Even in minimal extensions of the Standard Model (SM), such as the dark photon (A') or millicharged particles (mCPs), these predictions can carry significant theoretical uncertainty, particularly at proton beam experiments where hadronic production mechanisms are not well constrained. Such uncertainties propagate directly into signal acceptance estimates and ultimately affect the sensitivity reach of ongoing and future searches.

A key feature of A' and mCP models is that their interaction with SM particles is through electromagnetism. As a result, their production amplitudes are closely related to the amplitude for emission of an off-shell SM photon. This insight was used by [3] to develop a data-driven dark photon search strategy at the LHCb experiment. Here we extend this idea to fixed-target experiments in the context of both dark photons and millicharges. We will show that fully

differential distributions of virtual photons, which are accessible experimentally through dilepton measurements, encode precisely the information needed to predict the rate and kinematics of A' and mCP production.

In order to make use of the theoretical connection between A' /mCP production and off-shell photon emission in a realistic setting, we will employ a machine learning approach to model the dilepton kinematic distributions. Normalizing flows (NFs) provide a flexible way to achieve this and they offer fast sampling from the resulting model. By training a conditional NF on dilepton data, we construct a fully differential, data-driven Monte Carlo (MC) generator that can be used to simulate both dark photon and millicharged particle signals. This allows fixed-target experiments to anchor their signal modeling directly to measured SM processes.

The rest of this paper is organized as follows. In Sec. II we derive a simple relationship between the differential cross-sections for the production of A' or mCPs and dileptons. This result is the foundation of our method, enabling us to relate the dark photon or mCP rates and kinematics to observables in Sec. III. There we illustrate how to leverage this relationship in the context of a mock dataset tuned to reproduce published results from the NA60 fixed-target experiment. We also describe a normalizing flow architecture that can be used to model the observed kinematic distribution of dilepton pairs. This produces a data-driven Monte Carlo generator for A' or mCP kinematics. We discuss further applications of our approach and conclude in Sec. IV. Appendix A contains details of

*Contact author: nblinov@yorku.ca

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.

training and hyperparameter studies for the normalizing flow models.


II. PRODUCTION MECHANISMS

In order to effectively search for dark sector particles, one needs to identify the set of production channels relevant to each experiment, evaluate the kinematics of potential signal events, and compute the expected signal yield as a function of model parameters. At lepton beam experiments, dominant production modes of the dark photon include bremsstrahlung $\ell N \rightarrow \ell N + A'$, annihilation $e^+e^- \rightarrow A'$ and inverse Compton-like processes $\gamma e^- \rightarrow e^- A'$ [4–7]. The scattering matrix elements and kinematic distributions that follow from them are perturbatively calculable, even for bremsstrahlung, which involves scattering off the nucleus—the relevant hadronic form factor is known. Millicharged particles are produced through the same electromagnetic processes: replacing the final-state A' with $\tilde{\chi}\chi$ yields the corresponding mCP channels, whose rates can be computed with similar theoretical precision.

Hadronic beam experiments introduce a qualitatively different challenge. Dark photons and mCPs can be produced through numerous channels, including rare meson decays, hadronic bremsstrahlung, and mixing with vector mesons [8]. Some of these are straightforward to model, while others have significant uncertainties. Signal predictions typically rely on a combination of existing Monte Carlo generators, such as Pythia [9], or custom codes [10] designed for modeling low-energy hadronic interactions. These tools may not be validated in the low transverse momentum (p_T), forward region relevant for fixed-target searches. Additionally, some processes, such as coherent proton bremsstrahlung, are not included in most simulations and must be added by hand [11–15]. These calculations contain hadronic inputs, i.e., form factors, that are not currently well-constrained by data, and therefore lead to significant systematic uncertainties in signal predictions and data interpretation. Below we develop a method to calibrate dark photon and mCP production rates and kinematics *in situ* using dilepton measurements in the relevant detector, thereby bypassing many of the difficult modeling problems described above.

We imagine that an experiment has measured the kinematic distribution of off-shell SM photons through observations of, e.g., $\gamma^* \rightarrow \ell^+ \ell^-$. We want to leverage this measurement to predict the distribution and yield of A' or mCPs as a function of their coupling and mass. To do this, we will find a relationship between the γ^* and A' or mCP distributions.

First, we consider generic production of $\gamma^* \rightarrow \ell^+ \ell^-$. At the relatively low center-of-mass energies of interest for fixed-target experiments (well below the electroweak scale), the amplitude can be written schematically as



$$\mathcal{M}_{\gamma^*} = \text{A} \text{---} \gamma^* = \frac{e}{q^2} \mathcal{A}_\mu(q) \bar{u}_\ell \gamma^\mu v_{\bar{\ell}} \quad (1)$$

where $\mathcal{A}_\mu$ is part of the amplitude that describes the production of γ^* with (off-shell) momentum q ; u_ℓ and $v_{\bar{\ell}}$ are Dirac spinors corresponding to the outgoing leptons. The initial and final state can involve other particles, which are not shown. Using this parametrization we can compute the distribution $d\sigma_\gamma$, differential in γ^* mass $m_\gamma^2 \equiv q^2 = (p_{\ell^+} + p_{\ell^-})^2 = m_{\ell\ell}^2$, and three-momentum $\mathbf{q}$, with the result

$$\frac{d\sigma_\gamma}{dm_\gamma^2 d^3q} = \mathcal{F} \frac{\alpha}{48\pi^4 m_\gamma^2 q^0} \sqrt{1 - \frac{4m_\ell^2}{m_\gamma^2}} \left(1 + \frac{2m_\ell^2}{m_\gamma^2}\right) \times \int d\Phi \left[\frac{q^\mu q^\nu}{m_\gamma^2} - g^{\mu\nu} \right] \mathcal{A}_\mu \mathcal{A}_\nu^* \quad (2)$$

where $\alpha = e^2/4\pi$, $\mathcal{F}$ is the usual cross-section flux factor (which depends on the specific initial state), $q^0 = \sqrt{m_\gamma^2 + \mathbf{q}^2}$ is the off-shell photon energy, and $\int d\Phi$ is the Lorentz-invariant phase space integral over any final state particles except for $\ell^\pm$. To obtain the above distribution we inserted

$$1 = \int dm_\gamma^2 d^4q \delta^4(q - p_{\ell^+} - p_{\ell^-}) \delta(m_\gamma^2 - q^2) \theta(q^0) \quad (3)$$

and carried out the phase space integrals over $\ell^\pm$. We will perform similar calculations for the emission of A' and mCPs, and relate the results to Eq. (2).

A. Dark photon

Dark photons interact with the SM through a kinetic mixing with the photon [16,17].¹ Upon diagonalization, one finds that electromagnetic currents J_{EM}^μ couple to the A' just as they do to the photon, but with a strength suppressed by the dimensionless mixing parameter ϵ ,

$$\mathcal{L} \supset e(A_\mu + \epsilon A'_\mu) J_{\text{EM}}^\mu. \quad (4)$$

Using this interaction we compute the emission of a massive on-shell dark photon with polarization ϵ^μ ,



$$\mathcal{M}_{A'} = \text{A} \text{---} A' = \epsilon \mathcal{A}_\mu(q) u_\ell \epsilon^\mu(q). \quad (5)$$

¹We work far below the scale of electroweak symmetry breaking and do not consider the mixing with the hypercharge gauge boson.

Crucially, we assume that the amplitude $\mathcal{A}$ is the same as for γ^* ; we expect this to be a sound approximation since weak neutral currents would lead to corrections only of $\mathcal{O}(m_{A'}^2/m_Z^2)$.

Squaring Eq. (5) and summing over A' polarizations, we find a simple relationship between the dark photon differential cross-section $d\sigma_{A'}$ and $d\sigma_\gamma$,

$$\frac{d\sigma_{A'}}{d^3q} = \varepsilon^2 \left[\frac{\alpha}{3\pi} \sqrt{1 - \frac{4m_\ell^2}{m_{A'}^2}} \left(1 + \frac{2m_\ell^2}{m_{A'}^2} \right) \right]^{-1} \times m_{A'}^2 \frac{d\sigma_\gamma}{dm_\gamma^2 d^3q} \Big|_{m_\gamma^2 = m_{A'}^2}. \quad (6)$$

An analogous formula holds for decay rates with the replacement $\sigma_{A'} \rightarrow \Gamma_{A'}$, $\sigma_\gamma \rightarrow \Gamma$, where Γ is the decay width of a particle decaying into final states containing γ^* . For example, one can easily check that this relation holds between specific processes such as $\pi^0 \rightarrow A'\gamma$ and $\pi^0 \rightarrow \gamma(\gamma^* \rightarrow e^+e^-)$. Similar formulas appear in hadronic bremsstrahlung literature, relating cross-sections for virtual photon emission and dilepton production [18,19].

At proton beam experiments A' production is typically taken to be through three main channels: π^0 decay, η decay, and proton bremsstrahlung for $m_{A'} \lesssim 2$ GeV, see, e.g., [10,20,21]. Putting aside questions about the production kinematics of the parent π^0 and η mesons, their decay kinematics into A' are fixed, and the partial width can be accurately modeled using hadronic form factors determined in “single Dalitz” processes $\pi^0, \eta \rightarrow \gamma(\gamma^* \rightarrow e^+e^-)$.² Notably, the form factor measurements are at timelike momentum transfers, as required for decay rate calculations. Repurposing these SM measurements to estimate A' yield provides a simple application of Eq. (6) [23]. As an example, the uncertainties in recent determinations of the π^0 transition form factor [24] are $\sim 15\%$, and even better for η [25]. These are more than adequate for beyond Standard Model applications at the moment.

In contrast, models of proton bremsstrahlung require the introduction of form factors evaluated at q^2 outside of the observed range [12]. This leads to predictions that depend on unknown hadronic parameters. A common approach to approximate these form factors is to assume they take a dipole form with a single parameter $\Lambda \sim \text{GeV}$, a typical hadronic scale. Uncertainties on observables are then estimated by varying Λ . Our goal is to avoid these issues by using data and Eq. (6) to predict dark photon emission rate and kinematics as a function of its mass and coupling.

²These form factors are important for the dispersive determinations of $g-2$ [22].

B. Millicharge production

Millicharged particles χ couple to familiar matter via the electromagnetic interaction,

$$\mathcal{L} \supset e Q_\chi A_\mu \bar{\chi} \gamma^\mu \chi. \quad (7)$$

Such particles have been introduced as a component of the dark matter [26–28], to solve experimental anomalies [29,30], and as a generic feature of dark sector theories with massless kinetically mixed vectors [31]. A variety of fixed-target experiments have searched (or have results that have been reinterpreted) for mCPs [32–35]; many searches using existing or future facilities have been proposed, see, e.g., [33,36–38]. Importantly, mCP production inherits the same hadronic uncertainties as the dark photon, because in both cases the underlying process is the emission of an off-shell photon from the target. This makes mCPs a natural second application of the data-driven strategy developed in this work.

Much of the quantitative dark photon discussion also applies to mCPs, except now the parallels to dilepton production are even more straightforward. Specifically, Eq. (2) applies directly with the replacements $m_\ell \rightarrow m_\chi$ and $\alpha \rightarrow Q_\chi^2 \alpha$. We therefore find the following relation:

$$\frac{d\sigma_{\chi\chi}}{dm_\chi^2 d^3q} = Q_\chi^2 \frac{\sqrt{1 - \frac{4m_\ell^2}{m_\chi^2}} \left(1 + \frac{2m_\ell^2}{m_\chi^2} \right)}{\sqrt{1 - \frac{4m_\ell^2}{m_\gamma^2}} \left(1 + \frac{2m_\ell^2}{m_\gamma^2} \right)} \frac{d\sigma_\gamma}{dm_\gamma^2 d^3q}. \quad (8)$$

When applied to production from meson decays, e.g., $\pi^0 \rightarrow \gamma\bar{\chi}\chi$ and $J/\psi \rightarrow \bar{\chi}\chi$, the phase-space integrated version of Eq. (8) matches previous results, see, e.g., [36].

III. LEARNING FROM DATA

The main aim of this paper is to provide a recipe for extracting the distributions Eqs. (6) and (8) from data. Since we do not have access to real event-level data we first construct a mock dataset in Sec. III A. We will then relate the normalization of the A' or mCP rate (i.e., the total cross section) to the observed dilepton counts in Sec. III B. Finally, we will construct a fully differential generator for A' or mCP kinematics by training a conditional normalizing flow model on the observed kinematic distributions in Sec. III C.

For concreteness we work with NA60 results that provide detailed dimuon spectra, differential in invariant mass and transverse momentum [25,39], over a mass range relevant for dark sector searches. This experiment used a 400 GeV proton beam and a composite target. For simplicity we focus on the $L = 0.01$ interaction length lead subtarget of NA60, for which the most information is available. Similar measurements of e^+e^- pairs were made by NA45 [40]. In both cases the dilepton spectra were well




FIG. 1. Comparison of our tuned mock dataset (obtained by reweighing Pythia 8.3 samples) against dimuon invariant mass (left) and vector meson transverse momentum distributions observed by NA60 from their Pb target. In the right panel, the colored bands indicate 1σ ranges of the NA60 p_T spectrum fit parameters.

fit by “hadronic cocktail” models, consisting of two- and three-body decays of various light mesons. Both NA60 and NA45 observed forward dilepton production, with angular acceptances similar to future beam dump searches like SHiP [41,42] and DarkQuest [43], and we will assume that these future setups will be able to collect similar data sets. DarkQuest is an upgrade of SeaQuest, an experiment already designed to measure forward dimuon production. In contrast, SHiP is being designed with an active muon veto [44], so a dimuon measurement may require a short dedicated run without it.

A. Data and background

Our first step is to produce a mock dataset using a hadronic cocktail similar to NA60. To do this, we simulate pp collisions in Pythia 8.3 [9] and extract the hadrons used in NA60 fits. Specifically, the largest contributions to the dimuon signal come from two-body decays $\eta \rightarrow \mu^+\mu^-$, $\rho \rightarrow \mu^+\mu^-$, $\omega \rightarrow \mu^+\mu^-$, and three-body decays $\eta \rightarrow \gamma\mu^+\mu^-$, $\omega \rightarrow \pi^0\mu^+\mu^-$, and $\eta' \rightarrow \gamma\mu^+\mu^-$. We require the muons to have at least 3 GeV of energy, an approximate minimum needed for $\mu^\pm$ to pass through the NA60 hadron absorber. Each event corresponding to meson m is assigned a weight w_m according to measured branching fraction [45], total production cross-section σ_m , and detection efficiency $\mathcal{E}$ reported by NA60 [39],

$$w_m = N_{\text{POT}} \text{BR}(m \rightarrow \mu^+\mu^- X) L \frac{\sigma_m}{\sigma_{\text{inel}}} \mathcal{E}(m_{\ell\ell}^{(m)}, p_T^{(m)}), \quad (9)$$

where σ_{inel} is the total inelastic cross section in Pb that determines the nuclear interaction length. Unfortunately, we do not have complete information to precisely normalize our mock dataset: NA60 did not report the exact number of protons on target (POT) collected, N_{POT} , due to a detector malfunction, or the efficiency over the entire

dilepton invariant mass range. From [46] we surmise that $N_{\text{POT}} \sim 10^{13}$; the efficiency for three specific meson decays is provided in [39], which we interpolate to construct an approximate efficiency function $\mathcal{E}(m_{\ell\ell}^{(m)}, p_T^{(m)})$.

The weights w_m obtained via this procedure are not exact, but provide a rough approximation of the observations; to improve this further, we vary the meson contributions to best fit the published invariant mass spectrum [39]. The resulting weights are different from w_m by $\mathcal{O}(1)$ factors, which is easily attributed to lack of knowledge of the exact POT and acceptance. The resulting fit in terms of invariant mass is shown in the left panel of Fig. 1. Note that the observed spectrum is smeared by the 30 MeV detector resolution [46]. As a further check of the kinematics of our mock data, we have reproduced the $m_{\mu\mu}$ lineshapes for the three-body decays given in [25], and the meson p_T distribution (shown in the right panel of Fig. 1).³ We therefore conclude that our mock data is a reasonable representation of the experimental results.

The NA60 dimuon dataset includes a combinatorial background contribution at the level of 10%. These correspond to chance crossings of $\mu^+\mu^-$ tracks from unrelated meson decays that could not be removed through analysis selections. Since we will use event-level data, it will be important to study how the presence of these backgrounds affects the inference of the signal distribution. We therefore also produce a sample of combinatorial background as follows. Using the same Pythia 8.3 setup as above, we randomly pair opposite sign muons from $\pi^\pm$ and $K^\pm$ decays, weighing each event by the probability of the parent mesons to decay inside the target.

³To fully validate the kinematics of our mock data, we would also need to compare dimuon energy distributions (and their correlations with p_T), which have not been published.

Open charm production ($\bar{c}c$ with subsequent hadronization into D mesons and their decay into muon-containing final states) also provides a dimuon background, unrelated to γ^* production. Such events can be included in the same way as the combinatorial background; however, for simplicity we introduce a $m_{\gamma\gamma} < 1.1$ GeV cut (based on fits in [25]) to exclude these events when training our normalizing flow models below.

B. Production rate

We now apply Eqs. (6) and (8) to obtain dark photon/ m CP production rates and kinematics. Experimental selections like angular and energy cuts, detector efficiencies, and resolution effects sculpt the fundamental distribution of γ^* , leading to the shape in Fig. 1. These effects can be unfolded on an event-by-event basis [47,48]. However, proper unfolding requires a detailed detector simulation, which we do not have. To model the result of this unfolding we simply remove the detector acceptance factor from the weights described above. Resolution effects can also be unfolded, but we do not pursue this here for two reasons [49]. First, it would introduce uncertainties that would need to be tracked throughout the analysis. Second, resolution smearing smooths out sharp features in the invariant mass distribution, which simplifies training of the normalizing flow model below. Note that we need to exclude the contribution of $\eta \rightarrow \mu\mu$ from our dataset, since this process is not mediated by a single virtual photon [50]. The rate for this process can be extracted from data itself, via the $\eta \rightarrow \gamma\mu\mu$ contribution, and so the weight of the $m_{\mu\mu} \approx m_\eta$ bins can be reduced accordingly. Since we have access to the truth-level data, we simply omit the $\eta \rightarrow \mu\mu$ events in what follows.

The result of this simple unfolding procedure is a sample drawn from the distribution $d\sigma_\gamma/dm_\gamma^2 d^3q$, which determines the cross sections and distributions of dark photons and millicharged particles via Eqs. (6) and (8).

1. Dark photon

We compute the total A' cross section by binning our mock data in m_γ and evaluating $d\sigma_\gamma/dm_\gamma^2$ in each bin using the weights described above. The resulting yield per proton on target is shown in Fig. 2. Thus a yield prediction can be made *in situ*, using the detector and target configuration specific to a given experiment.

Alongside our (mock) data-driven prediction of the A' yield we show a theoretical calculation of proton-nucleus bremsstrahlung $pN \rightarrow A' + X$, where X is an arbitrary hadronic final state. This calculation used the QRAv2 prescription from [14], form-factor parameters from [13], and we normalized the rate using the nonsingle diffractive proton-nucleus cross section given in [51]. The form factor used did not include a contribution from the ϕ meson, so the corresponding spike in the invariant mass spectrum at




FIG. 2. Data-driven estimate of A' yield (black line) for a NA60-like beam and target configuration as a function of $m_{A'}$. The mock data includes only meson decays, and not bremsstrahlung (red curve), as the latter contribution was not required to describe the NA60 data [25]. The data-driven result accounts for the effects of detector resolution. The red band corresponds to variations of a proton form-factor parameter, in the bremsstrahlung calculation by a factor of 2. Note that both curves correspond to the full phase space without any angular cuts.

$m_{A'} \approx 1$ GeV is absent; more sophisticated models [15,52,53] do feature such an enhancement. Note that bremsstrahlung was not one of the processes we used to construct our mock dataset, so it is interesting that the yields are comparable in regions where vector meson decays dominate.⁴ The similarity in rates, however, hides striking differences in detailed kinematic distributions, which we will comment on later. Only real data can disentangle different contributions to vector production at these invariant masses.

The result in Fig. 2 is a “raw” yield estimate, without any angle cuts (i.e., corresponding to the full phase space of A') or requirements on the detection of the dark photon. For realistic sensitivity predictions one also needs to compute the acceptance for the specific signal to be searched for. In general, this acceptance function is not the same one that determines the SM dimuon yield ($\mathcal{E}$ above). For example, if the goal is to look for displaced decays of the A' the new acceptance function will depend on the location of A' decay vertex, directions of its decay products, etc—quantities that all depend on the energy and momenta of the emitted A' . A realistic calculation of the acceptance is therefore only possible if we can sample A' kinematics at each point of its parameter space (i.e., for different values $m_{A'}$). We show an example of such an acceptance calculation using our trained NFs in Appendix B.

⁴Vector mesons kinematics are generated in Pythia 8.3 through hadronization of a partonic initial state, rather than coherent bremsstrahlung.

Finally, we note that if the observed distribution $d\sigma_\gamma/dm_\gamma^2 d^3q$ is contaminated by combinatorial backgrounds, the predicted signal yield can be corrected by subtracting the combinatorial yield. This can be inferred from data by studying same-sign lepton pairs, whose production rate is similar to the opposite-sign dileptons that contribute to the background. This is precisely the procedure followed by NA60 to obtain the background-subtracted spectrum shown in Fig. 1.

2. Millicharged particles

We obtain the mCP yield by summing the aforementioned weights for events with $m_\gamma > 2m_\chi$, multiplied by the prefactor in Eq. (8). The result is shown in Fig. 3. Here we note that experimental selections (such as the focus on dimuon final states, or cuts on $m_{\mu\mu}$) can exclude important production mechanisms, so the (mock) data-driven prediction may not be sufficient. We illustrate this point in Fig. 3 by including production from $\pi^0 \rightarrow \gamma\bar{\chi}\chi$ and $J/\psi \rightarrow \bar{\chi}\chi$, following [36]. These production modes were normalized using pPb π^0 and J/ψ yields from [54] and [55,56], respectively. Similar to the dark photon case, one can also consider mCP production from hadronic bremsstrahlung $pN \rightarrow X + \gamma^*(\bar{\chi}\chi)$. We follow the prescription of [15] (via [37]) to estimate this yield; this result, shown as the red band in Fig. 3, is subject to the same hadronic uncertainties as the dark photon case.

C. Kinematic distribution

1. Dark photon

A measurement of the dilepton invariant mass requires knowledge of the energy q^0 and momentum of the virtual photon $\mathbf{q}$. Thus such a dataset should provide a fully differential estimate of $d\sigma_\gamma/dm_\gamma^2 d^3q$. To make predictions for a dark photon model with a specific $m_{A'}$, however, we are interested in this distribution with the restriction $m_\gamma = m_{A'}$. In other words, we want to know the conditional probability density

$$p(\mathbf{q}|m_\gamma = m_{A'}) \propto \left. \frac{d\sigma_\gamma}{dm_\gamma^2 d^3q} \right|_{m_\gamma=m_{A'}}. \quad (10)$$

We will further assume azimuthal symmetry to reduce the dimensionality of our distribution; this means that in practice we will work with q_T , the transverse momentum (with respect to the beam axis), and $E_\gamma = q^0$, the energy of virtual photon/dark photon, instead of the full three-momentum $\mathbf{q}$.

There are many ways to perform this density estimation. A simple nonparametric approach is to bin the data in all of the kinematic variables. Kinematic samples can be drawn using the accept/reject method on the histogram at the desired $m_{A'}$ bin. While simple, such a method suffers from




FIG. 3. Data-driven estimate of mCP (χ) yield (black line) for an NA60-like beam and target configuration as a function of m_χ . The mock data includes only meson decays needed to describe the NA60 result and not bremsstrahlung (red curve), which may not account for all relevant production channels in a given mass window. For example, for NA60, contributions from $\pi^0 \rightarrow \gamma\bar{\chi}\chi$ and $J/\psi \rightarrow \bar{\chi}\chi$ (shown as colored lines), are not included due to experimental selections.

the usual trade-offs between resolution and statistical noise stemming from bin size choice; additionally the underlying smooth distribution is modeled by a piecewise-constant function. More sophisticated methods of density estimation such as kernel density estimation and Gaussian mixture models alleviate these concerns but may have other disadvantages, such as difficulty representing power-law distribution tails (cf. right panel of Fig. 1), inability to model conditional distributions in a simple way, or cumbersome sampling procedure. We avoid these issues by using a normalizing flow (NF) [57,58] to model the conditional distribution in Eq. (10).

Normalizing flows map simple base distributions (e.g., a multivariate Gaussian) to complex target distributions by applying a series of continuous, invertible, and differentiable transformations [59–62]. This enables estimation of complex probability densities, evaluation, and efficient sampling of the resulting models. Flow models work by finding a transformation f that relates a data point $\mathbf{x}$ to a latent random variable $\mathbf{z}_0$, which is distributed according to the base distribution. The mapping f is represented by a composition of K invertible transformations f_i ,

$$\mathbf{x} = f(\mathbf{z}_0) = (f_K \circ \dots \circ f_1)(\mathbf{z}_0). \quad (11)$$

Each f_i is parametrized by a set of tunable parameters, typically implemented using neural networks. Change-of-variables and chain rule formulas allow one to relate the corresponding data $[p(\mathbf{x})]$ and base $[p_0(\mathbf{z}_0)]$ distributions as

$$\ln p(\mathbf{x}) = \ln p_0(\mathbf{z}_0) - \sum_{i=1}^K \ln \left| \det \left(\frac{\partial \mathbf{z}_i}{\partial \mathbf{z}_{i-1}} \right) \right|, \quad (12)$$

where $\mathbf{z}_i \equiv (f_i \circ \dots \circ f_1)(\mathbf{z}_0)$. The functions f_k must be carefully chosen to be both bijective, and to enable efficient evaluation of the Jacobian determinants in Eq. (12). Training the NF model involves maximizing this log-likelihood over the dataset by varying the parameters that specify f_i . Once the flow is trained, we can sample $\mathbf{z}_0 \sim p(\mathbf{z}_0)$ and then apply the flow transformation to obtain $\mathbf{x} = f(\mathbf{z}_0)$, a new sample from the model distribution. In conditional NF models, f_i additionally depend on a context variable $\mathbf{c}$, which can be discrete or continuous [63]. This allows the flow to learn conditional distributions $p(\mathbf{x}|\mathbf{c})$ and sample from the trained model for a given value of $\mathbf{c}$.

Normalizing flows have enjoyed a multitude of applications in particle physics [45,64]. Closest in spirit to our use case are works that used these models to accelerate Monte Carlo integration/sampling of differential cross sections or detector simulation [65–73]. There a normalizing flow is trained on samples from a theoretically known (either analytically or via a simulation) distribution to construct an approximation that is efficient to sample from. We instead want to train a NF on the observed (and unfolded) kinematic distributions to enable the generation of new samples, conditioned on one of the variables.

In practice we use the following standard architecture for the flow, implemented with the help of NORMFLOWS [74] and PyTorch [75]. We use an autoregressive rational quadratic spline flow with $K = 8$ flow steps, with a linear transformation of the latent variables between each step [76]. Each flow transformation f_i is a monotonic rational quadratic spline with six bins. The parameters of these splines are modeled by autoregressive neural networks [77] with two hidden layers of 32 features each. We assume azimuthal symmetry of the detector, so the flow acts on pairs $\mathbf{x} = (q_T, E_\gamma)$, with the corresponding invariant masses m_γ supplied as the conditioning parameters $\mathbf{c}$. In Appendix A we show that broad ranges of model and optimizer hyperparameters provide similar performance. For the results presented in the main text, we train on the unweighted mock dataset consisting of $\sim 3 \times 10^5$ events. As expected, the NF performs best in regions with high statistics; performance degrades in sparsely populated mass bins (i.e., high m_γ). This can be alleviated by weighing the high- m_γ data more in training, either using a weighted likelihood [65] or a weighted training data sampler.

Note that our mock dataset is larger than the data collected by NA60 (cf. event counts in the left panel of Fig. 1). NA60’s POT was relatively modest compared to those targeted by future dark sector searches: for example, DarkQuest aims to collect over 10^{18} POT, so its dimuon dataset can be significantly larger. Another important difference between NA60 and many dark sector experiments is target thickness:




FIG. 4. Kullback-Leibler divergence between normalizing flow model trained on mock data and the data itself, as a function of the virtual photon invariant mass m_γ . The green (yellow) bands show 68% (95%) confidence limits for KL divergence between different realizations of the data. The black dotted lines indicate the 95% uncertainty bands due to finite sample size drawn from the NF model. The NF model is compatible with the data distribution up to statistical uncertainties, which are more severe at larger invariant masses because of smaller sample sizes.

the latter experiments often use thick targets and dilepton production can occur at different positions, which may also be added to the training dataset. The distribution of these positions can be important for evaluating the acceptance of long-lived particles.

To validate the resulting models we compute the Kullback-Leibler divergence between the NF model, p_m , and the mock data distribution, p_d [78,79],

$$\text{KL}(p_m \| p_d) = \int d\mathbf{x} p_m(\mathbf{x}) \ln \left(\frac{p_d(\mathbf{x})}{p_m(\mathbf{x})} \right), \quad (13)$$

where $\mathbf{x} = (q_T, E_\gamma)$ and both distributions are conditioned on m_γ .⁵ This quantity vanishes when the two distributions are identical. While the normalizing flow architecture allows for an easy evaluation of p_m , we only have samples from p_d . We thus estimate KL using a discrete approximation, following [80,81], with finite samples drawn from the model at different $m_{A'}$ and a subset of the mock data restricted to invariant masses in a small window around $m_{A'}$. This metric is shown in Fig. 4 as the solid black line. In order to understand what constitutes a “good” value of this metric, we compare this “model vs data” result to Kullback-Leibler (KL) between different draws from the mock data set. In the limit of infinite sample size this should vanish. However, depending on the invariant mass region considered, the amount of data varies significantly for a fixed

⁵The conditional data distribution p_d is approximated using the histogram approach described above, by binning the data in m_γ according to the detector resolution.




FIG. 5. Comparison of the normalized q_T (left column) and $E_{A'} = E_\gamma$ (right column) dark photon distributions sampled from two different NF models and the mock signal dataset for different $m_{A'} = m_\gamma$ values. The filled histograms show the training data (background contribution) in blue (orange) in each invariant mass bin. The model labelled “NF” (light gray) was trained on a signal + combinatorial background sample, while “NF (no bg)” (dark gray) was trained on a signal-only sample. In each row the data is a subset of the full dataset with virtual photon invariant masses within 30 MeV of the corresponding $m_{A'}$. Both NF models successfully interpolate the distributions to arbitrary $m_{A'}$. The model trained on the realistic sample that includes the background provides a good description of the signal-only distribution, with differences only becoming apparent in the tails.

sample size, leading to statistical fluctuations in KL. Reassuringly, the “data vs data” divergence (green dashed line) is close to zero across the invariant mass range, while “model vs data” is within one standard deviation from this, with larger excursions occurring at values of invariant mass where there were fewer training samples (c.f. Fig. 1). In Appendix A we illustrate how KL improved with larger training sets across the invariant mass range.

As described above, one of the main advantages of normalizing flows for density estimation is the ability to easily sample from the trained model. Because during training m_γ is provided as a continuous context, the model is able to interpolate between different values of m_γ . In Fig. 5 we show distributions of kinematic samples drawn from the model at three different values of $m_{A'} = m_\gamma$. While the dataset contains a range of $m_\gamma \in [2m_\mu, 1.1 \text{ GeV}]$, the specific $m_{A'}$ values shown in that figure are not present as discrete data points in the training set. However, because the underlying distributions are smooth, they should be close to the observed distribution of events averaged over a narrow range centered on each invariant mass value. We see that this is indeed the case, across the entire relevant invariant mass range.

We conclude that the NF model can be well trained with reasonable dataset sizes, and offers a faithful approximation to $p(q_T, E_\gamma | m_{A'})$. This enables fast generation of A' kinematics with arbitrary mass (as long as they lie in the range probed by the data), without introducing theoretical

models for specific production channels. This data-driven generator can then be used to evaluate efficiency functions for, e.g., long-lived A' searches. This quantity, together with the overall A' yield normalization discussed in Sec. III B, can be used to accurately evaluate the sensitivity of such observations.

2. Millicharged particles

In contrast to the dark photon case, a wide range $m_\gamma \geq 2m_\chi$ of virtual photon invariant masses contributes to mCP production. Normalizing flows can be used to construct a data-driven MC here as well. Several approaches are possible. The most straightforward is to train a different flow for each m_χ hypothesis, enforcing $m_\gamma \geq 2m_\chi$ at the training set level. Another possibility is to use the entire m_γ range available in the data, and then perform rejection sampling on $m_\gamma = m_{\chi\chi}$ at generation time. Only one unconditional flow is needed in this case. Finally, one can train the same conditional flow as for the dark photon, and use it in conjunction with sampling from the one-dimensional invariant mass distribution. In all of these cases, once a viable γ^* sample is generated, it can be easily decayed into $\bar{\chi}\chi$, which can be used for signal acceptance calculations. We have validated the latter approach and obtained comparable agreement to the dark photon case (see Fig. 5) between NF-generated and truth-level (i.e., obtained directly from our mock data) $\bar{\chi}\chi$ distributions.

IV. DISCUSSION

We have shown that fixed-target dark photon or milli-charged particle production can be inferred directly from measured dilepton spectra. The method makes use of a special property of these scenarios: all interactions are electromagnetic in nature. This leads to simple relationships, Eqs. (6) and (8), between the A' or mCP cross section and the observed virtual photon distribution, enabling both overall rate predictions and fully differential kinematic sampling without detailed hadronic modeling.

We further demonstrated that conditional normalizing flows can be used to model the observed kinematic distribution. A key feature of these techniques is the ease of sampling from the resulting model. This enables realistic and fast simulation of dark sector signal events, which is necessary to accurately determine signal acceptance.

A number of questions and possible extensions follow from this strategy. First, the dark photon is only one example of a gauge extension of the SM. Can the same technique be applied to other anomaly-free vector particles, like $B-L$ and L_i-L_j vector bosons? Certainly the situation is more complicated compared to the A' model, since those vector bosons do not just couple to the EM current. However, in certain cases, such a coupling is generated at loop level, as in the $L_\mu-L_\tau$ model. Despite the loop suppression, this induced kinetic mixing is experimentally important [82] and so our method would apply directly to this case. This approach can also be applied to other currents if their SM signatures can be isolated experimentally in analogy to dilepton events.

Another interesting application of fully differential kinematic distributions would be to disentangle the roles of different production mechanisms of vector mesons and beyond-SM particles. In Figs. 2 and 3 we illustrated how a data-driven estimate of dark sector particle production can be compared to a theoretical calculation of bremsstrahlung. While there are uncertain inputs (the proton form-factor) to these calculations, the overall rate appears to be consistent with “data.” However, if we compare the kinematic distributions in Fig. 6 we find striking differences (the same observation holds for mCP bremsstrahlung). We emphasize that ultimately Fig. 6 is a comparison between Pythia 8.3, which produces vector mesons via hadronization, and coherent emission of A' or $\gamma^*(\bar{\chi}\chi)$ from a proton. Only real observations can quantify the relative contributions of these mechanisms in different regions of phase space. In particular, Fig. 6 suggests that vector energy distributions could be a powerful discriminator between hadronic bremsstrahlung and hadronization for both A' and mCP models. Curiously hadronic bremsstrahlung of dilepton pairs also has an uncertain history; it was initially invoked to explain excess dilepton production in high-energy fixed-target experiments at low invariant masses [40]. It turned out that the excess was due to underestimated η yield, with the bremsstrahlung contribution being highly subdominant [83]. At low, $\sim$ GeV,




FIG. 6. Predicted A' phase space distribution in the $(E_{A'}, q_T)$ plane for $m_{A'} = 0.7$ GeV, sampled from the NF model trained on mock data (signal and combinatorial background, black lines) or from bremsstrahlung (red line). While the q_T distributions are similar, the energy distributions are distinct. The inner (outer) contours indicate 1σ (2σ) containment regions.

beam energies, however, hadronic bremsstrahlung is significant, see, e.g., [84].

Our proposed analysis neglects several important experimental subtleties. First, our desired SM signal, off-shell photons traced by dilepton pairs, is subject to backgrounds. For example, there are combinatorial backgrounds due to chance overlaps of independent muon tracks coming from, e.g., charged meson decays. Much of this background can be eliminated on an event-by-event basis using tracking and tight vertex selections. At NA60, the residual combinatorial background was about 10% of the total dataset [25]. After these selections, it is difficult to reject specific events as being combinatorial, but the overall rate of the backgrounds and the shape of their distribution can be estimated from data. Another source of dilepton events are Bethe-Heitler reactions $\gamma Z \rightarrow Z \ell^+ \ell^-$ and $\eta \rightarrow \ell^+ \ell^-$, where ℓ^+ and ℓ^- emerge from different vertices in both cases. These events do not contribute to $\gamma^* \rightarrow \ell^+ \ell^-$ and therefore should be excluded from the analysis. The contribution of these events to the desired signal can be reduced with kinematic cuts, or by adjusting the weight of the corresponding phase space. The upshot is that realistic event samples are likely not pure samples of virtual photons. Another source of uncertainty is introduced during the unfolding of dilepton acceptance, which is associated with knowledge of the detector, Monte Carlo or data statistics. All of these ambiguities should translate into uncertainties of the kinematic samples and their relative weights. These can be propagated to the A' model in the Bayesian normalizing

flow framework [85]; in this architecture the flow parameters are modeled as distributions, rather than fixed numbers. Model uncertainties from these multiple sources can then be assessed by averaging predictions over multiple realizations of the flow parameters [86].

In the future it will be important to assess the utility of our proposed method in the context of actual future facilities, accounting for their target thickness, detector geometry, and uncertainties as outlined above.

This work is made possible by the following open-source software: NORMFLOWS [74], PyTorch [75], SCIKIT-LEARN [87], CORNER.PY [88], VEGAS [89], NumPy [90], SciPy [91] and Matplotlib [92].

ACKNOWLEDGMENTS

We thank Yoni Kahn, Shirley Li, Kevin Kelly, and Sean Tulin for useful discussions. We are especially grateful to Maxim Pospelov for discussions that initiated this project and for helpful feedback on the manuscript. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). Our numerical calculations were enabled in part by support provided by Compute Ontario, Calcul Québec, and the Digital Research Alliance of Canada [93].

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: NORMALIZING FLOW TRAINING

In this appendix we discuss the details of our normalizing flow model training and hyperparameter selection. Before training, our dataset was log transformed and standardized by subtracting the mean and dividing by the standard deviation. We trained our model using the ADAM algorithm [94] for 500 epochs in batches of 1024 points, a learning rate of 10^{-3} and an optimizer weight decay of 10^{-5} . In each epoch 5×10^4 data points are randomly drawn from our data set for the model to train on.

Normalizing flow model hyperparameters were chosen by performing a systematic scan over the number of flow steps, hidden layers, and features. The mean KL divergence was then computed for each set of hyperparameters over 100 different seed values, so that model-generated data was varied for each KL calculation. KL divergence calculations used $k = 2$ nearest neighbors for the density approximation. Since a range of values were found to provide a low mean KL divergence between model-generated points and the data, the final hyperparameter values were chosen to minimize architectural complexity and the number of tunable parameters within the NF transformation. Hyperparameter scans were performed on data without

the combinatorial background. All plots in this appendix used data with background contribution.

The dataset was 10^6 points generated from 10^8 Pythia 8.3 events, with 10^5 points used for plots within the main text, while the full data set was used for Appendix figures. The dataset was split with 80% used for training and 20% for validation. Since past works have shown test data should be at least the same size as training data [95], another set of 10^6 points was used for testing model results and KL divergence calculations. The training loss was taken as the mean negative log likelihood over all batches of the training data per epoch. The validation loss was calculated once per epoch as the negative log likelihood of the full set of the




FIG. 7. Typical training and validation losses as a function of training epoch for the normalizing flow model.




FIG. 8. Kullback-Leibler divergence between the normalizing flow model and the data for different m_γ values as a function of the number of points used to train the model. As in Fig. 4, the green (yellow) bands show 68% (95%) confidence limits for KL divergence between different realizations of the data. Unlike Fig. 4, the black dotted lines now indicate the 1σ uncertainty bands due to initialization variance, i.e., different initial weights of the NF model prior to training.




FIG. 9. Long-lived particle detector acceptance as a function of decay length for three representative mass points. The acceptance is estimated using Eq. (B1) for a mock detector described in the text. The solid and dashed lines correspond to A' samples generated using NF models trained on mock data with or without the combinatorial background contribution.

validation data. A typical set of loss curves as a function of training epoch is shown in Fig. 7, showing good convergence and absence of overfitting (which would be indicated by an increase in only the validation loss).

Training results in normalizing flow models have been shown to depend on neural network initialization [81,96]. To quantify this, the model was trained 50 times with random seeds corresponding to different starting model weights. The standard deviation of the mean KL divergence over these runs was 0.002, 0.006, and 0.012 for $m_\gamma = 0.3$, 0.7 and 1.1 GeV, respectively, highlighting initialization variance as another source of uncertainty in our model. Figure 8 shows the KL divergence for different m_γ values as a function of the amount of data used in model training, averaged over different model weight initializations. As the KL divergence between the model and the data is lower for denser m_γ regions, bias from our approximation method for the mock data distribution becomes more visible [97], causing the model KL divergence to fall outside of the data KL divergence range. However results still showed that initialization variance decreased as the amount of training points increased. In different architectures, wider networks have been shown to reduce bias [98].

APPENDIX B: IMPACT OF COMBINATORIAL BACKGROUND ON ACCEPTANCE

We have noted that realistic dilepton datasets will include a contribution from events that do not originate from $\gamma^* \rightarrow \mu^+ \mu^-$ decays. These events arise due to chance overlaps of opposite sign muons from unrelated meson decays, or from “open-charm” processes where $c\bar{c}$ production is followed by hadronization and muonic decay of the resulting D mesons. In this appendix we quantify the impact of this data contamination on estimates of dark photon acceptance in a prototypical long-lived particle search.

We consider a mock detector, with geometric acceptance in the range $\theta_{\min} = 35 \text{ mrad} < \theta_{A'} < \theta_{\max} = 120 \text{ mrad}$

relative to the beam axis, and a $L = 5 \text{ m}$ shield (motivated by the NA60 setup). We then use our NF models to sample dark photon momenta and compute the acceptance $\mathcal{E}$

$$\mathcal{E} = \langle \Theta(\theta_{A'} - \theta_{\min}) \Theta(\theta_{\max} - \theta_{A'}) e^{-L/(\gamma v \tau_{A'})} \rangle \quad (\text{B1})$$

where the angle brackets represent an average over the samples. We repeat this using NF trained on data with and without combinatorial background. The acceptance for three representative mass points is shown in Fig. 9 as function of $c\tau_{A'}$. We observe good qualitative agreement between the two NF models.

The acceptance in Eq. (B1) can be exponentially sensitive to the kinematic distribution. We therefore compute the relative difference $2(\mathcal{E}_{\text{w bkg}} - \mathcal{E}_{\text{no bkg}})/(\mathcal{E}_{\text{w bkg}} + \mathcal{E}_{\text{no bkg}})$ between background and no-background models across a broad range of parameter space in Fig. 9. As expected, the acceptance in the long-lived regime is only mildly sensitive to the energy distribution, leading to




FIG. 10. Relative difference in acceptance computed using NF models trained on data with and without combinatorial background in the $m_{A'}\text{-}c\tau_{A'}$ plane.

relative differences of less than 10% across all masses considered. Shorter lifetimes accentuate the differences between predicted distributions as shown in Fig. 10, but the signal yield is also exponentially falling in that parameter space. An assessment of the impact of this

effect would require a more detailed analysis using the total expected yield in a given experiment. Moreover, if a systematic shift in the acceptance can be computed using mock data as illustrated here, it can be corrected for in evaluating the true data-driven acceptance prediction.

-
- [1] J. Beacham *et al.*, *J. Phys. G* **47**, 010501 (2020).
 - [2] B. Batell, N. Blinov, C. Hearty, and R. McGehee, in *Snowmass 2021* (2022), [arXiv:2207.06905](#).
 - [3] P. Ilten, Y. Soreq, J. Thaler, M. Williams, and W. Xue, *Phys. Rev. Lett.* **116**, 251803 (2016).
 - [4] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, *Phys. Rev. D* **80**, 075018 (2009).
 - [5] E. Nardi, C. D. R. Carvajal, A. Ghoshal, D. Meloni, and M. Raggi, *Phys. Rev. D* **97**, 095004 (2018).
 - [6] L. Marsicano, M. Battaglieri, M. Bondi', C. D. R. Carvajal, A. Celentano, M. De Napoli, R. De Vita, E. Nardi, M. Raggi, and P. Valente, *Phys. Rev. D* **98**, 015031 (2018).
 - [7] A. Celentano, L. Darmé, L. Marsicano, and E. Nardi, *Phys. Rev. D* **102**, 075026 (2020).
 - [8] Y. Kyselov and M. Ovchinnikov, *Phys. Rev. D* **111**, 015030 (2025).
 - [9] C. Bierlich *et al.*, *SciPost Phys. Codebases* **2022**, 8 (2022).
 - [10] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz, *Phys. Rev. D* **95**, 035006 (2017).
 - [11] J. Blümlein and J. Brunner, *Phys. Lett. B* **731**, 320 (2014).
 - [12] S. Foroughi-Abari and A. Ritz, *Phys. Rev. D* **105**, 095045 (2022).
 - [13] D. Gorbunov and E. Kriukova, *J. High Energy Phys.* **01** (2024) 058.
 - [14] S. Foroughi-Abari, P. Reimitz, and A. Ritz, *Phys. Rev. D* **112**, 015030 (2025).
 - [15] F. Kling, P. Reimitz, and A. Ritz, *Phys. Rev. D* **113**, 015008 (2026).
 - [16] B. Holdom, *Phys. Lett.* **166B**, 196 (1986).
 - [17] F. del Aguila, G. D. Coughlan, and M. Quiros, *Nucl. Phys.* **B307**, 633 (1988); **B312**, 751(E) (1989).
 - [18] P. Lichard, *Phys. Rev. D* **49**, 5812 (1994).
 - [19] P. Lichard, *Phys. Rev. D* **51**, 6017 (1995).
 - [20] A. Berlin, S. Gori, P. Schuster, and N. Toro, *Phys. Rev. D* **98**, 035011 (2018).
 - [21] J. M. Berryman, A. de Gouvea, P. J. Fox, B. J. Kayser, K. J. Kelly, and J. L. Raaf, *J. High Energy Phys.* **02** (2020) 174.
 - [22] I. Danilkin, C. F. Redmer, and M. Vanderhaeghen, *Prog. Part. Nucl. Phys.* **107**, 20 (2019).
 - [23] B. Batell, M. Pospelov, and A. Ritz, *Phys. Rev. D* **80**, 095024 (2009).
 - [24] C. Lazzeroni *et al.* (NA62 Collaboration), *Phys. Lett. B* **768**, 38 (2017).
 - [25] R. Arnaldi *et al.* (NA60 Collaboration), *Phys. Lett. B* **757**, 437 (2016).
 - [26] S. L. Dubovsky, D. S. Gorbunov, and G. I. Rubtsov, *JETP Lett.* **79**, 1 (2004).
 - [27] S. D. McDermott, H.-B. Yu, and K. M. Zurek, *Phys. Rev. D* **83**, 063509 (2011).
 - [28] C. Dvorkin, T. Lin, and K. Schutz, *Phys. Rev. D* **99**, 115009 (2019); **105**, 119901(E) (2022).
 - [29] A. Berlin, D. Hooper, G. Krnjaic, and S. D. McDermott, *Phys. Rev. Lett.* **121**, 011102 (2018).
 - [30] J. B. Muñoz and A. Loeb, *Nature (London)* **557**, 684 (2018).
 - [31] B. Holdom, *Phys. Lett. B* **178**, 65 (1986).
 - [32] A. A. Prinz *et al.*, *Phys. Rev. Lett.* **81**, 1175 (1998).
 - [33] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, *Phys. Rev. Lett.* **122**, 071801 (2019).
 - [34] R. Acciarri *et al.* (ArgoNeuT Collaboration), *Phys. Rev. Lett.* **124**, 131801 (2020).
 - [35] L. Barak *et al.* (SENSEI Collaboration), *Phys. Rev. Lett.* **133**, 071801 (2024).
 - [36] K. J. Kelly and Y.-D. Tsai, *Phys. Rev. D* **100**, 015043 (2019).
 - [37] S. N. Gninenko, D. V. Kirpichnikov, and N. V. Krasnikov, *Phys. Rev. D* **100**, 035003 (2019).
 - [38] S. Choi *et al.*, [arXiv:2007.06329](#).
 - [39] R. Arnaldi *et al.* (NA60 Collaboration), *Eur. Phys. J. C* **79**, 443 (2019).
 - [40] G. Agakichiev *et al.*, *Eur. Phys. J. C* **4**, 231 (1998).
 - [41] O. Aberle *et al.* (SHiP Collaboration), BDF/SHiP at the ECN3 high-intensity beam facility, Technical Report, CERN, Geneva, 2022.
 - [42] R. Albanese *et al.* (SHiP Collaboration), BDF/SHiP at the ECN3 high-intensity beam facility, Technical Report, CERN, Geneva, 2023.
 - [43] A. Apyan *et al.*, in *Snowmass 2021* (2022), [arXiv:2203.08322](#).
 - [44] A. Akmete *et al.* (SHiP Collaboration), *J. Instrum.* **12**, P05011 (2017).
 - [45] S. Navas *et al.* (Particle Data Group), *Phys. Rev. D* **110**, 030001 (2024).
 - [46] A. Uras, Low mass dimuon production in p-A collisions at 400 GeV/c with the NA60 detector, Ph.D. thesis, Cagliari University (2011).
 - [47] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, and J. Thaler, *Phys. Rev. Lett.* **124**, 182001 (2020).
 - [48] F. Canelli *et al.*, *Eur. Phys. J. C* **86**, 106 (2026).
 - [49] G. Cowan, *Statistical Data Analysis* (Oxford University Press, 1998).
 - [50] L. Gan, B. Kubis, E. Passemar, and S. Tulin, *Phys. Rep.* **945**, 1 (2022).
 - [51] N. Blinov, E. Kowalczyk, and M. Wynne, *J. High Energy Phys.* **02** (2022) 036.

- [52] D. Gorbunov and E. Kriukova, *J. High Energy Phys.* **02** (2025) 018.
- [53] D. Gorbunov and E. Kriukova, *Phys. Part. Nucl.* **56**, 506 (2025).
- [54] J. Cincheza *et al.*, *Nucl. Phys.* **B158**, 280 (1979).
- [55] B. Alessandro *et al.* (NA50 Collaboration), *Eur. Phys. J. C* **48**, 329 (2006).
- [56] R. Arnaldi *et al.* (NA60 Collaboration), *Phys. Lett. B* **706**, 263 (2012).
- [57] I. Kobzyev, S. J. Prince, and M. A. Brubaker, *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3964 (2020), <https://ieeexplore.ieee.org/document/9089305>.
- [58] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, *J. Mach. Learn. Res.* **22**, 1 (2021), [arXiv:1912.02762](https://arxiv.org/abs/1912.02762).
- [59] E. G. Tabak and E. Vanden-Eijnden, *Commun. Math. Sci.* **8**, 217 (2010).
- [60] E. G. Tabak and C. V. Turner, *Commun. Pure Appl. Math.* **66**, 145 (2013).
- [61] L. Dinh, D. Krueger, and Y. Bengio, [arXiv:1410.8516](https://arxiv.org/abs/1410.8516).
- [62] D. Rezende and S. Mohamed, in *International Conference on Machine Learning* (PMLR, Lille, France, 2015), pp. 1530–1538.
- [63] C. Winkler, D. Worrall, E. Hoozeboom, and M. Welling, [arXiv:1912.00042](https://arxiv.org/abs/1912.00042).
- [64] HEP ML Community, A Living Review of Machine Learning for Particle Physics, <https://iml-wg.github.io/HEPML-LivingReview/>.
- [65] B. Stienen and R. Verheyen, *SciPost Phys.* **10**, 038 (2021).
- [66] E. Bothmann, T. Janßen, M. Knobbe, T. Schmale, and S. Schumann, *SciPost Phys.* **8**, 069 (2020).
- [67] C. Gao, J. Isaacson, and C. Krause, *Mach. Learn. Sci. Tech.* **1**, 045023 (2020).
- [68] C. Gao, S. Höche, J. Isaacson, C. Krause, and H. Schulz, *Phys. Rev. D* **101**, 076002 (2020).
- [69] T. Heimel, R. Winterhalder, A. Butter, J. Isaacson, C. Krause, F. Maltoni, O. Mattelaer, and T. Plehn, *SciPost Phys.* **15**, 141 (2023).
- [70] N. Deutschmann and N. Götz, *J. High Energy Phys.* **03** (2024) 083.
- [71] E. Bothmann, T. Janßen, M. Knobbe, B. Schmitzer, and F. Sinz, [arXiv:2506.18987](https://arxiv.org/abs/2506.18987).
- [72] C. Krause and D. Shih, *Phys. Rev. D* **107**, 113003 (2023).
- [73] C. Krause and D. Shih, *Phys. Rev. D* **107**, 113004 (2023).
- [74] V. Stimper, D. Liu, A. Campbell, V. Berenz, L. Ryll, B. Schölkopf, and J. M. Hernández-Lobato, *J. Open Source Software* **8**, 5361 (2023).
- [75] J. Ansel *et al.*, in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, Volume 2, ASPLOS '24 (Association for Computing Machinery, New York, USA, 2024), pp. 929–947.
- [76] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, in *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)* (2019), [arXiv:1906.04032](https://arxiv.org/abs/1906.04032).
- [77] M. Germain, K. Gregor, I. Murray, and H. Larochelle, in *International Conference on Machine Learning* (PMLR, Lille, France, 2015), pp. 881–889.
- [78] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (Springer, New York, 2002).
- [79] C. M. Bishop and N. M. Nasrabadi, *Pattern Recognition and Machine Learning* (Springer, New York, 2006), Vol. 4.
- [80] Q. Wang, S. R. Kulkarni, and S. Verdú, in *2006 IEEE International Symposium on Information Theory* (IEEE, Seattle, WA, USA, 2006), pp. 242–246.
- [81] Y. Kahn, S. Roy, J. Shelton, and V. Tiki, in *NeurIPS Workshop* (NeurIPS, Vancouver, BC, Canada, 2024).
- [82] N. Blinov, P. J. Fox, K. J. Kelly, R. Plestid, and T. Zhou, [arXiv:2511.09619](https://arxiv.org/abs/2511.09619).
- [83] T. Åkesson *et al.*, *Z. Phys. C* **68**, 47 (1995).
- [84] G. Agakishiev *et al.* (HADES Collaboration), *Phys. Lett. B* **690**, 118 (2010).
- [85] B. L. Trippe and R. E. Turner, [arXiv:1802.04908](https://arxiv.org/abs/1802.04908).
- [86] C. Bierlich, P. Ilten, T. Menzo, S. Mrenna, M. Szewc, M. K. Wilkinson, A. Youssef, and J. Zupan, *SciPost Phys.* **17**, 045 (2024).
- [87] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *J. Mach. Learn. Res.* **12**, 2825 (2011), [arXiv:1201.0490](https://arxiv.org/abs/1201.0490).
- [88] D. Foreman-Mackey, *J. Open Source Software* **1**, 24 (2016).
- [89] G. P. Lepage, *J. Comput. Phys.* **439**, 110386 (2021).
- [90] C. R. Harris *et al.*, *Nature (London)* **585**, 357 (2020).
- [91] P. Virtanen *et al.* (SciPy 1.0 Contributors), *Nat. Methods* **17**, 261 (2020).
- [92] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [93] <http://alliancecan.ca>.
- [94] D. P. Kingma, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [95] A. Cocco, M. Letizia, H. Reyes-González, and R. Torre, *Symmetry* **16**, 942 (2024).
- [96] C. M. Bishop and H. Bishop, *Deep Learning—Foundations and Concepts* (Springer, New York, 2024).
- [97] P. Zhao and L. Lai, *IEEE Trans. Inf. Theory* **66**, 7787 (2020).
- [98] B. Neal, S. Mittal, A. Baratin, V. Tanti, M. Scicluna, S. Lacoste-Julien, and I. Mitliagkas, [arXiv:1810.08591](https://arxiv.org/abs/1810.08591).