LUNAR: A Monte Carlo generator for bound-nucleon decay in liquid argon

J. A. Nowak*

School of Physics and Astronomy, Lancaster University, Lancaster LA1 4YB, United Kingdom

(Received 1 July 2026; revised 17 August 2026; accepted 11 September 2026; published 27 October 2026)

The search for nucleon decay in liquid-argon time-projection chambers requires a quantitative description of how the bound nuclear environment reshapes the decay-product kinematics. We present LUNAR, a fast, openly available Monte Carlo generator dedicated to two-body decays of protons and neutrons bound in argon-40, the target nucleus of the DUNE far detector. The parent nucleon is drawn from a selectable nuclear ground state (ten momentum distributions ranging from mean-field Fermi gases to argon spectral functions) and bound off the mass shell by one of three removal-energy prescriptions, including the momentum-dependent optical potential of Juszczak *et al.* The two-body decay is performed off shell and boosted to the laboratory frame, and the daughter meson is then propagated out of the nucleus by a semiclassical intranuclear cascade with an optional formation zone for the freshly produced meson. We use the generator to separate the distinct roles of Fermi motion and binding in shaping the observable meson spectrum, to quantify final-state interactions channel by channel, and to translate present Super-Kamiokande limits into expected DUNE event yields for the full set of standard decay modes. Final-state interactions leave the supersymmetry-favored $p \to K^+\bar{\nu}$ signal, and the neutral-kaon modes, essentially intact while roughly halving the pion, η, and antikaon rates, an effect that dominates over the $\sim$10% spread induced by the choice of nuclear model. For the pion channels, the modeling uncertainty of the cascade itself exceeds that spread. The code is released to the community as a lightweight, extensible tool for signal efficiency and systematics studies.

DOI: 10.1103/mw7s-jl1c

I. INTRODUCTION

The conservation of the baryon number is an accidental symmetry of the Standard Model, not supported by any gauge principle, and essentially every attempt to embed the Standard Model into a more unified framework predicts that it is broken. Grand unified theories are the canonical example: by placing quarks and leptons in common multiplets, they introduce interactions that convert one into the other, so that the proton (and the neutron bound inside a nucleus) becomes unstable. The minimal $SU(5)$ theory of Georgi and Glashow [1], rooted in the quark-lepton unification of Pati and Salam [2], mediates decay through dimension-six operators [3,4] and singles out $p \to e^+\pi^0$ as the leading mode. Supersymmetric grand unification changes this picture qualitatively: dimension-five operators generated by the exchange of colored Higgsinos dominate, and, because they favor a strange quark in the final state, the preferred channel becomes $p \to K^+\bar{\nu}$ [5,6]. Which mode is dominant is, therefore, a direct diagnostic of the underlying unification scheme, and a broad program of searches across many channels is needed to discriminate between them [7,8]. Baryon-number violation need not stop at single-nucleon decay; dinucleon transitions and neutron-antineutron oscillation probe complementary operators [9], and a flexible simulation tool should anticipate them.

Liquid-argon time-projection chambers, and, in particular, the far detector of the Deep Underground Neutrino Experiment (DUNE), are exceptionally well matched to the supersymmetry-favored channel [10]. A kaon produced in the decay of a free proton at rest would be monochromatic and, at the relevant momentum, would stop and decay inside the active volume, leaving a short, heavily ionizing track with a distinctive decay topology. This combination of calorimetric and tracking information makes $p \to K^+\bar{\nu}$ the golden liquid-argon channel. The same detector retains sensitivity to the full complement of standard two-body modes: $e^+\pi^0$, $\mu^+\pi^0$, $\bar{\nu}\pi$, the e^+ and μ^+ companions of the η, and the neutral- and charged-kaon channels, each of which already carries a dedicated experimental limit (Sec. II D).

The free-nucleon case, however, is never realized in a detector. The decaying nucleon is bound, so it carries momentum and sits below its free mass by a removal energy that depends on the nuclear shell it occupies. Both effects act on the observable: Fermi motion smears the

*Contact author: j.nowak@lancaster.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International 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.

otherwise sharp daughter spectrum, while binding pushes the decaying system off the mass shell, lowering the available energy and, for the heavier final states, forbidding a fraction of decays outright. The decay meson must then traverse the residual nucleus, where it can scatter, charge exchange, or be absorbed before it ever reaches the detector. An honest estimate of the signal efficiency, of the reconstructed meson spectrum, and of the accompanying hadronic activity, therefore, demands a controlled model of the nuclear initial state and of the meson's passage out of the nucleus.

These nuclear effects are, in principle, within reach of the general-purpose neutrino event generators GENIE [11], NuWro [12], and GiBUU [13], whose nuclear models and intranuclear cascades were built for the same argon target. Those frameworks, however, are engineered around the neutrino-nucleus interaction vertex; using them to study nucleon decay means working against their intended design, and their breadth comes at the cost of weight and opacity. Recent work has begun to quantify these effects specifically for proton decay: a GiBUU transport calculation [14] for $p \to e^+\pi^0$ in water Cherenkov detectors and a survey by Barbu, Parvu, and Lazanu [15] of Fermi motion and final-state interactions (FSIs) for $p \to \pi^+\bar{\nu}$ and $p \to K^+\bar{\nu}$ across liquid argon, xenon, and water. Both, however, either repurpose a heavy general-purpose transport code or work outside an event-generator framework altogether. There is, consequently, a gap for a small, transparent, dedicated tool, in the spirit of compact single-purpose generators such as MARLEY [16], that an analyzer can read in an afternoon, run in seconds, and repurpose to a new channel or nuclear model by editing a single header. LUNAR (the Lancaster University nucleon-decay argon generator) is built to fill that gap. It samples a bound nucleon from a menu of nuclear models, performs the off-shell two-body decay, transports the daughter meson through a semiclassical cascade, and writes a complete final-state record, for any of the standard proton and neutron channels.

In this paper, we describe the physics content of the generator (Sec. II) and use it to address three questions of direct relevance to a liquid-argon nucleon-decay search. First, we disentangle the separate effects of Fermi motion and nuclear binding on the observable meson spectrum, showing that the momentum model alone fixes its *width* while both ingredients move its *central value*, the momentum model, by the larger amount (Sec. III). Second, we quantify final-state interactions channel by channel and demonstrate that they are negligible for the strange modes but a leading, strongly mode-dependent correction for the pion, η, and antikaon channels. Third, we fold these ingredients into expected DUNE event yields for every standard mode, anchored to current Super-Kamiokande limits. We then discuss how the generator can be used and extended by the community (Secs. IV and V) before concluding (Sec. VI).

II. THE GENERATOR

LUNAR is written as a header-only C++17 physics library driven by a small set of command-line executables. It has no external dependencies beyond the standard library; the ROOT framework is used only by the downstream analysis macros, not by the generator itself. A single run samples the bound-nucleon ground state, performs the off-shell decay, propagates the products through the nucleus, and streams the final state to a plain-text file (Fig. 1). The physics is organized into independent modules: nuclear momentum and binding, off-shell kinematics, the spectral functions, the hadron-nucleon cross sections, and the cascade. Each of these can be modified or replaced without touching the others. This modularity is deliberate: it is what makes the extensions of Sec. V a matter of local edits rather than redesign.

A. Initial-state nucleon momentum

The argon-40 ground state is described in the local-density approximation. The nucleon density follows a two-parameter Fermi (Woods-Saxon) form [17]:

$$\rho(r) = \frac{\rho_0}{1 + \exp[(r - C)/C_1]}, \qquad \rho_0 = 0.176\ \text{fm}^{-3}, \quad (1)$$

with $C = 3.530$ fm and $C_1 = 0.541$ fm. Treating each nucleon species as a locally cold Fermi gas, the position-dependent Fermi momentum is

$$k_{\text{F}}^q(r) = \hbar c \left(3\pi^2 \frac{N_q}{A} \rho(r)\right)^{1/3}, \quad (2)$$

evaluated separately for protons and neutrons with the argon isospin fractions $Z/A = 0.45$ and $N/A = 0.55$. For protons, the density-weighted average reproduces the

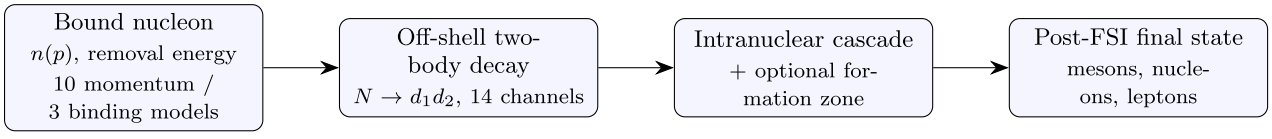


FIG. 1. The LUNAR generation pipeline. A bound nucleon is drawn from the chosen momentum and binding model, decayed off the mass shell into the selected channel, and its hadron is transported out of the nucleus by the intranuclear cascade (with an optional formation zone), yielding the complete post-FSI final state. Each stage is an independent, replaceable module.

TABLE I. Initial-state nucleon momentum models available in LUNAR. Parameter choices (high-momentum fractions, cutoffs, oscillator length) are summarized in Appendix A.

Key	Model	$n(p)$
POLY	Polynomial (toy)	$f(p)$ from config file
GFG	Global Fermi gas	$\propto p^2$, $p < k_F$
LFG	Local Fermi gas	Sphere $k_F(r)$, $r^2\rho(r)$-weighted
SRC	Short-range corr.	Fermi bulk $+1/p^4$ tail
SF	Spectral fn. (toy)	Local Fermi gas + tail
HOSM	HO shell model	$\sum_{nl}$ occ$\lvert\phi_{nl}(p)\rvert^2$
BR	Bodek-Ritchie	Fermi bulk + long $1/p^4$ tail
GAUSS	Gaussian	$\propto p^2 \exp(-p^2/2\sigma^2)$
CFG	Correlated FG	Sharp sphere + contact tail
Benhar	Tabulated $S(p,E)$	NuWro argon grid
Ankowski	Analytic $S(p,E)$	Shell mean field $+1/p^4$ tail

reference value $k_F^p = 217$ MeV/c, which also serves as the fixed Fermi momentum of the global-gas option; the corresponding neutron value is $k_F^n = 230$ MeV/c (Appendix A). Unless stated otherwise, the numbers quoted below refer to proton modes, and k_F means k_F^p.

The user selects the nucleon momentum distribution $n(p)$ from the ten physics-based models in Table I; an 11th option, a polynomial parametrization (POLY) whose coefficients are read from a configuration file, is a deliberately simple toy that can be tuned to prototype an arbitrary $n(p)$ for testing rather than to represent a specific nuclear model. The ten physics models span a deliberate range of physical assumptions. The global Fermi gas is a hard sphere with a sharp surface at k_F and no strength above it. The local-density variants soften that surface, populating low momenta because a decay can occur in the dilute nuclear periphery. The correlated models (short-range correlations, the Bodek-Ritchie gas [18], and the correlated Fermi gas) restore the high-momentum nucleons that a pure mean field omits, adding a $1/p^4$ tail above k_F. The harmonic-oscillator shell model replaces the sphere altogether by a sum of single-particle orbitals, producing genuine shell structure, while the Gaussian provides a smooth reference with no Fermi surface at all. Finally, the two spectral-function options supply their own joint momentum and removal-energy distribution (Sec. II C): benhar draws the nucleon directly from a tabulated argon spectral function, while ankowski builds an effective spectral function analytically from the argon shell structure and a correlated tail.

The resulting distributions are overlaid in Fig. 2, and a few observables are collected in Table II: the mean and root mean square of the nucleon momentum, the position of the distribution peak, and the fraction of nucleons carrying more than k_F. The latter is the cleanest discriminator between the model families. The global Fermi gas, by construction, has none; the local-density and shell models put a tenth to a fifth of the strength above k_F; and the correlated models and the spectral functions carry between a fifth and a third there. As we show in Sec. III, this high-momentum content is what controls the width of the observable daughter spectrum.

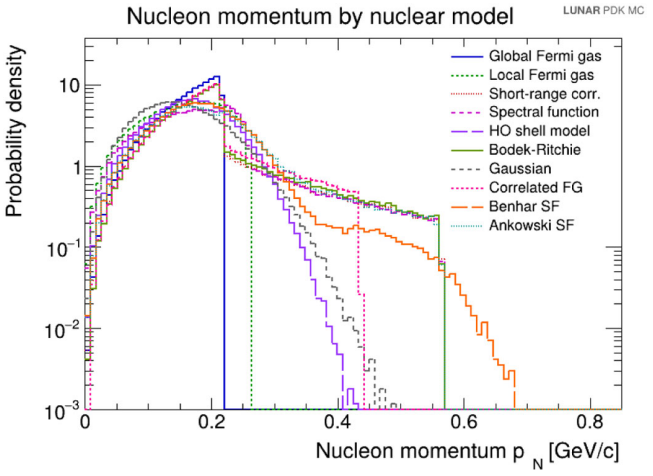


FIG. 2. Proton momentum distributions generated by LUNAR for the nuclear models in Table I (area-normalized, logarithmic vertical scale). The local-density treatment smears the Fermi-gas edge near k_F, while the short-range-correlation, Bodek-Ritchie, and correlated-Fermi-gas options extend well beyond it; the Gaussian falls off smoothly and the harmonic-oscillator shell model shows orbital structure in place of a sharp surface. Models are distinguished by line style as well as by color.

B. Binding and off-shell kinematics

A bound nucleon is not a free particle at rest: it is off the mass shell by an amount governed by the nuclear potential it feels. LUNAR offers three interchangeable prescriptions for the removal energy of a mean-field nucleon, selected at run time with the—binding flag:

TABLE II. Initial nucleon-momentum observables by model (2×10^5 $p \to K^+\bar{\nu}$ events, optical-potential binding). The high-momentum fraction $p_N > k_F^p = 0.217$ GeV/c rises from zero for the hard-edged global gas to roughly a third for the correlated and spectral-function models. The peak is the center of the most populated bin in Fig. 2.

Model	$\langle p_N \rangle$ [GeV/c]	rms [GeV/c]	Peak [GeV/c]	$p_N > k_F$ [%]
Global Fermi gas	0.163	0.042	0.208	0.0
Local Fermi gas	0.157	0.058	0.174	17.5
Short-range corr.	0.196	0.088	0.208	18.8
Spectral function	0.191	0.097	0.225	32.9
HO shell model	0.168	0.060	0.174	20.8
Bodek-Ritchie	0.199	0.090	0.208	20.7
Gaussian	0.155	0.065	0.149	17.1
Correlated FG	0.190	0.072	0.208	19.9
Benhar SF	0.189	0.080	0.174	30.4
Ankowski SF	0.200	0.094	0.174	35.4

(i) Potential. The momentum-dependent optical potential of Eq. (3) below (the default), which is the deepest and the only momentum-dependent choice;
(ii) Constant. A fixed average separation energy, 30 MeV by default;
(iii) Shell. The argon-40 single-particle levels, each Gaussian-smeared about its separation energy and weighted by occupancy.

Nucleons drawn from a short-range-correlated pair are not described by a mean field and always take the two-body recoil energy of Eq. (4) instead, whichever prescription is selected; the spectral-function models in Sec. II C supply their own joint momentum and removal-energy distribution and ignore the flag altogether. We describe the default in detail below, since it is the one used throughout Sec. III, and quantify the difference between the three prescriptions in Sec. III.

For the mean-field nucleons, the default takes the real part of the momentum-dependent optical potential of Ref. [17]:

$$V(k_{\mathrm{F}}, p) = -\frac{(ak_{\mathrm{F}})^2(k_{\mathrm{F}} + b)}{c^4 + d^3 k_{\mathrm{F}} + e^3 p^2/k_{\mathrm{F}} + p^4}, \tag{3}$$

with $a = 206$ MeV, $b = 582$ MeV, $c = -322$ MeV, $d = 422$ MeV, and $e = 289$ MeV. It is attractive, deepest for slow nucleons, and tends smoothly to zero at high momentum; at the argon Fermi momentum, it reaches $V(p = 0) = -59$ MeV, consistent with Ref. [17].

The potential fixes the energy of the bound nucleon N in the nuclear rest frame, $E_N = \sqrt{p^2 + M_N^2} + V(k_{\mathrm{F}}, p)$, which, in turn, defines an effective removal energy $E_{\mathrm{rem}} = M_N - E_N$ and an off-shell invariant mass $W^2 = E_N^2 - p^2$ that is the true energy budget of the decay. Here, p is always the nucleon momentum and N the bound nucleon, proton, or neutron. For the correlated nucleons, the removal energy is set instead by the recoil of the correlated partner:

$$E_{\mathrm{rem}} \simeq E_0 + \frac{p^2}{2M_N}, \tag{4}$$

where $E_0 = 20$ MeV is a fixed offset; the same E_0 sets the correlated ridge of the effective spectral function in Sec. II C. When the binding and Fermi motion drive W below the summed daughter mass, the decay is kinematically impossible and the nucleon is resampled.

Figure 3 shows the off-shell mass distribution for the kaon mode across the nuclear models, together with the forbidden fraction it implies. Binding and Fermi motion shift the whole distribution below the free mass M_p, and the correlated and spectral-function models develop a low-W tail fed by their deeply bound, high-momentum nucleons. That tail is the whole story for the forbidden fraction, and the right-hand panel quantifies it for three channels of increasing summed daughter mass. The mean-field models (global and local Fermi gas, harmonic-oscillator shell model, Gaussian, and the correlated Fermi gas, whose tail is cut at $2k_{\mathrm{F}}$) never reach below the threshold of any of the three channels and lose nothing at all. The models with an extended $1/p^4$ tail do: for $p \to K^+\bar{\nu}$ (threshold 0.494 GeV) they forbid 1.5%–1.8% of draws, rising to 5.5% for Bodek-Ritchie, whose tail extends to 1 GeV/c. Raising the threshold raises the loss monotonically: 2.2% (6.1% for Bodek-Ritchie) for $p \to e^+\eta$ at 0.548 GeV and 3.5%–4.1% (8.0%) for $p \to \mu^+\eta$ at 0.654 GeV, the heaviest two-body final state in the channel list. The effect is, therefore, small in absolute terms but entirely model dependent, and it is one more reason to quote nuclear-model spreads rather than

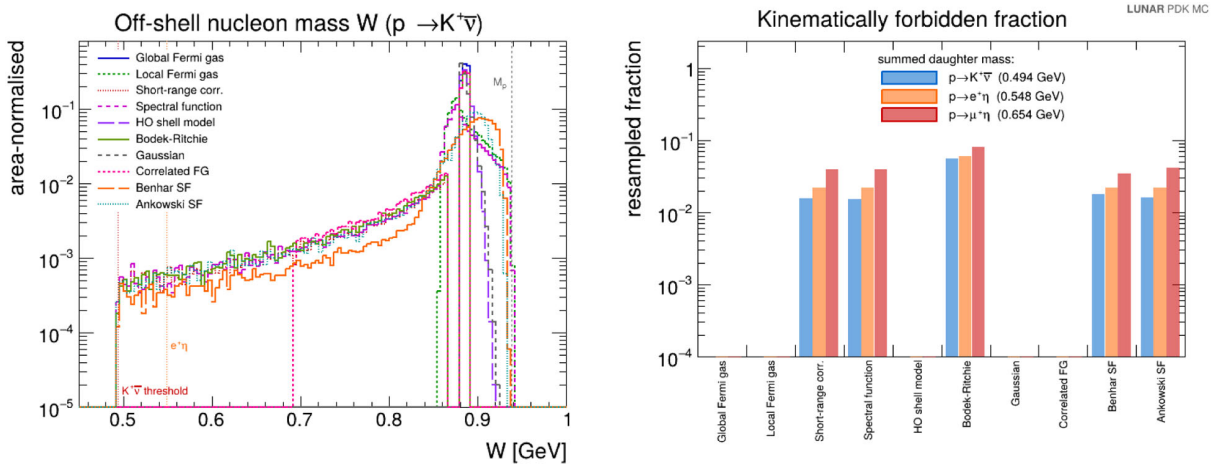


FIG. 3. Left: the off-shell nucleon invariant mass $W = \sqrt{(M_N - E_{\mathrm{rem}})^2 - p^2}$ for $p \to K^+\bar{\nu}$, by momentum model (area-normalized, logarithmic scale); the dashed line marks the free proton mass and the dotted lines the $K^+\bar{\nu}$ and $e^+\eta$ thresholds. Right: the resulting kinematically forbidden (resampled) fraction per model, for three channels of increasing summed daughter mass (logarithmic scale; the mean-field models forbid nothing at all and so have no visible bar). Binding pulls W well below M_p and the correlated tails feed the low-W region; the loss is confined to the models that populate it and grows with the threshold.

a single prediction. Because forbidden draws are resampled, the spectra shown throughout this paper are conditioned on the decay being kinematically allowed; read as an absolute in-medium rate, a channel would be suppressed by the corresponding few percent. We do not fold this into the yields of Sec. III C, where it is well inside the nuclear-model band.

C. Spectral functions

The benhar and ankowski options bypass the analytic momentum sampling entirely and instead draw the momentum and removal energy *jointly* from a spectral function $S(p,E)$: the momentum from the p^2-weighted marginal $\int S(p,E)dE$ and the removal energy from $S(p,\cdot)$ conditioned on it. Because $S(p,E)$ already encodes both the mean-field shell structure and the correlated tail in the full (p,E) plane, no separate high-momentum tail or binding model is layered on top. The two options differ in how $S(p,E)$ is obtained.

The benhar option reads a tabulated grid. The default grids are the Jefferson Lab E12-14-012 argon proton and neutron spectral functions [19,20] as distributed with NuWro [12], chosen automatically according to the parent nucleon. They are not hard-wired: the—sf-file option replaces the built-in choice with *any* tabulated $S(p,E)$ supplied by the user, so a different argon calculation, or a grid for another nucleus, can be used without touching the code. The expected format is the plain-text NuWro grid-fun2d layout: a three-line header giving the number of removal-energy and momentum points and the two axis ranges in MeV, followed by one block per momentum value listing the (E,S) pairs (Appendix B). Only relative values matter, and the sampler builds its marginal and conditional distributions from whatever grid it is given. Figure 4 displays the default grids themselves: a bright mean-field band at low momentum and low removal energy and a diffuse correlated ridge reaching to high p and high E.

The ankowski option instead constructs an *effective* Ankowski-Sobczyk spectral function [21] analytically, with no external grid. A fraction $f_{\rm corr}=0.20$ of nucleons populate a correlated $1/p^4$ tail carrying the two-nucleon removal energy $E_{\rm rem}=E_0+p^2/2M_N$; the remaining nucleons form a shell mean field, in which each argon orbital contributes its harmonic-oscillator momentum profile $|\phi_{nl}(p)|^2$ paired with a Gaussian-broadened separation energy from the $(e,e'p)$ shell table (Table I's HOSM profiles and the shell separation energies, combined into one joint distribution). The two parts are normalized so that the correlated strength is exactly $f_{\rm corr}$ in $\int 4\pi p^2 S dp dE$. The resulting $S(p,E)$, shown in Fig. 5, reproduces the same qualitative anatomy as the tabulated grid (a shell band stacked at the orbital separation energies plus a correlated ridge climbing with momentum) while remaining fully analytic and grid-free and gives a measurably broader nucleon-momentum distribution than the benhar grid (Table II).

Two features of this construction deserve comment. The value $f_{\rm corr}=0.20$ is not tuned here: it is the standard correlated strength of the effective spectral function of Ref. [21], and it agrees with the ∼20% of nucleons that many-body calculations place outside the mean field in medium and heavy nuclei [22] and with the short-range-correlated fraction extracted from electron scattering [23,24]. The sharp ridge that this produces in the (p,E) plane is not a physical feature of the nucleus but an artifact of the approximation: the correlated pair is treated as a nucleon recoiling against a single spectator at rest, so all of its removal-energy strength collapses onto the line $E_{\rm rem}=E_0+p^2/2M_N$. In a full calculation the center-of-mass motion of the pair and the excitation spectrum of the residual $A-2$ system smear that line into the broad ridge

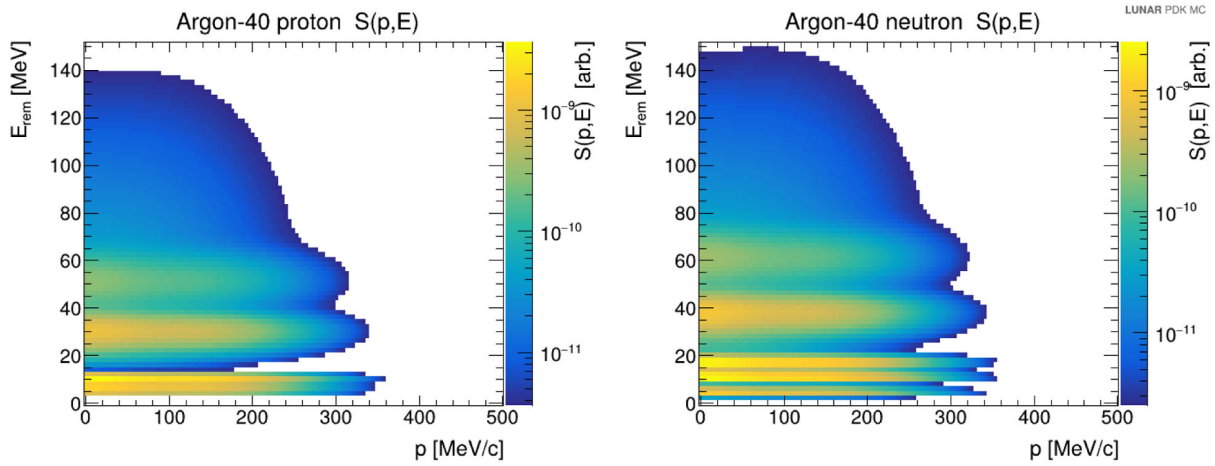


FIG. 4. Tabulated argon-40 spectral function $S(p,E)$ (logarithmic color scale) for the proton (left) and neutron (right) grids [19,20] that drive the benhar model. The bright low-(p,E) region is the mean-field shell structure; the diffuse high-momentum, high-removal-energy strength is the short-range-correlated tail.

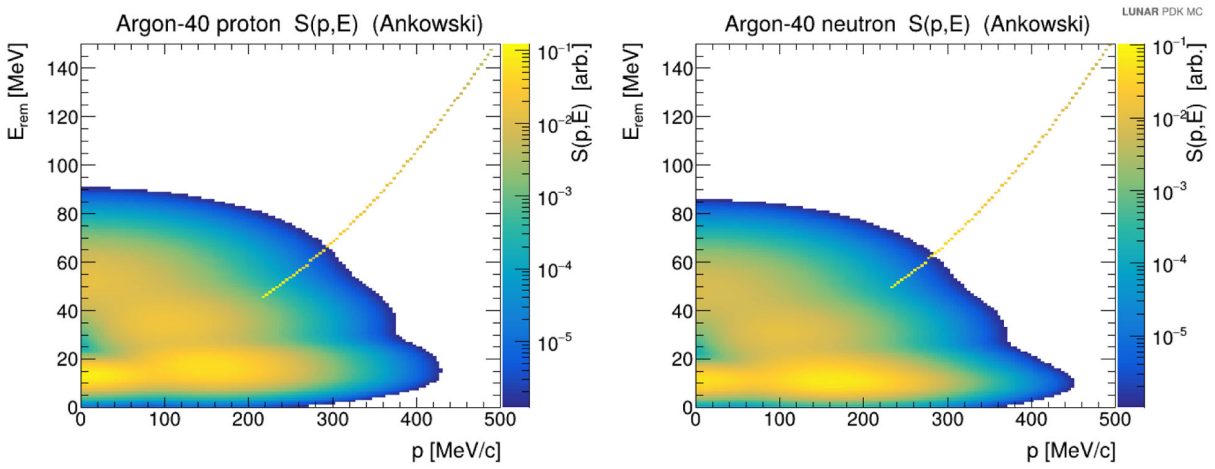


FIG. 5. Analytic Ankowski-Sobczyk effective spectral function $S(p, E)$ (logarithmic color scale) built by the ankowski model for the proton (left) and neutron (right), with no grid input. The horizontal band is the shell mean field at the orbital separation energies; the thin ridge climbing with momentum is the correlated $1/p^4$ tail at $E_{\rm rem} = E_0 + p^2/2M_N$ [Eq. (4)]. The neutron band reaches to higher momentum through its occupied $1f_{7/2}$ orbital. Compare with the tabulated grid in Fig. 4.

visible in the tabulated grid in Fig. 4. The two options are provided precisely so that this can be tested: the effective form is fast, analytic, and transparent, the tabulated grid carries the realistic width, and comparing them bounds what the idealization costs. As Sec. III shows, the resulting kaon spectra differ by only a few MeV$/c$ in mean and rms, because the observable integrates over the ridge.

D. Off-shell two-body decay and channels

Given the off-shell mass W, a generic two-body decay $N \to d_1 + d_2$ is generated relativistically. With the Källén function $\lambda(a, b, c) = a^2 + b^2 + c^2 - 2(ab + bc + ca)$, the daughter momentum in the nucleon rest frame is

$$p^* = \frac{\sqrt{\lambda(W^2, m_1^2, m_2^2)}}{2W}, \tag{5}$$

the two daughters are emitted back to back with an isotropic orientation, and both four-vectors are boosted to the laboratory frame with the nucleon velocity $\boldsymbol{\beta} = \mathbf{p}/E$. For the default $p \to K^+\bar{\nu}$ mode, the massless antineutrino reduces Eq. (5) to $p^* = (W^2 - m_K^2)/2W$. The nucleon masses are $M_p = 0.93827$ GeV and $M_n = 0.93957$ GeV.

The 14 supported channels are listed in Table III. They cover the standard supersymmetry- and GUT-motivated two-body modes for both bound protons and neutrons, and almost all of them carry a dedicated Super-Kamiokande search: the strange modes [25–27], the π^0 modes [28], the $\bar{\nu}\pi$ modes [29], the charged-antilepton-plus-meson modes [30], and the η modes [31]. The neutron channels reuse the same machinery with the neutron mass, Fermi momentum, isospin fraction, and shell table. Each mode produces a lepton-side daughter (an e^+, a μ^+, or an invisible $\bar{\nu}$) and a hadron-side meson whose laboratory momentum is the primary observable.

TABLE III. Nucleon-decay channels supported by LUNAR, selected with the—channel flag.

Key	Mode	Key	Mode
pToKnu	$p \to K^+\bar{\nu}$	nToEPim	$n \to e^+\pi^-$
pToEPi0	$p \to e^+\pi^0$	nToMuPim	$n \to \mu^+\pi^-$
pToMuPi0	$p \to \mu^+\pi^0$	nToNuPi0	$n \to \bar{\nu}\pi^0$
pToNuPip	$p \to \bar{\nu}\pi^+$	nToNuEta	$n \to \bar{\nu}\eta$
pToEEta	$p \to e^+\eta$	nToNuK0	$n \to \bar{\nu}K^0$
pToMuEta	$p \to \mu^+\eta$	nToEKm	$n \to e^+K^-$
pToEK0	$p \to e^+K^0$		
pToMuK0	$p \to \mu^+K^0$		

E. Final-state interactions

The decay meson is born inside the nucleus and must work its way out, and on the way it interacts with the spectator nucleons. To capture this, LUNAR transports the decay hadron and every secondary it creates through a semiclassical intranuclear cascade, switched on by default. The cascade follows the approach of NuWro's kaskada [12,32], itself descended from the classic intranuclear-cascade calculations [33,34] and the pion-transport simulations of Salcedo, Oset, and collaborators [35,36], and shares its philosophy with the cascade final-state models of modern neutrino generators [11].

The decay happens where the decaying nucleon is. For the local-density models, which sample a radius r from the $r^2\rho(r)$ weight of Eq. (1) in order to evaluate $k_{\rm F}(r)$, the cascade reuses *that* radius as the decay vertex, so that the momentum, the binding, and the vertex all refer to the same point of the nucleus; only the direction of the vertex is drawn afresh. The remaining models sample the nucleon

momentum from a global distribution and, therefore, define no radius; for them, the vertex is drawn independently from the same $r^2\rho(r)$ weight, which reproduces the correct marginal vertex distribution but not its correlation with the nucleon momentum. The distinction matters little in practice (the escaping-meson yields of Sec. III B move by less than a percentage point when the correlation is switched off), because the two treatments differ only in that correlation and not in the distribution of vertices itself.

The hadron is then advanced in short straight steps of 0.05 fm; at each step the local mean free path is $\lambda = 1/(\rho\sigma)$, with σ the hadron-nucleon total cross section resolved into its proton and neutron parts. When an interaction occurs, a target nucleon is drawn from a local Fermi sphere of radius $k_{\mathrm{F}}^{q}(r)$ [Eq. (2)] evaluated at the current position and for the species q of the struck nucleon. The spectators are, thus, always treated as a local Fermi gas, independently of the model chosen for the *decaying* nucleon: the momentum models in Table I describe the single nucleon that decays, including its short-range-correlated tail, while the medium the daughter meson traverses is the mean field. A subprocess (elastic scattering, charge exchange, absorption, or multimeson production) is chosen in proportion to its partial cross section, the final state is generated in the projectile-target rest frame, and Pauli blocking is imposed on every outgoing nucleon. Secondaries are cascaded until all particles have either escaped or been absorbed. Four-momentum and electric charge are conserved exactly at every vertex, and the leptons, which do not interact strongly, pass through untouched.

Quasielastic and charge-exchange scattering are forward peaked, $\mathrm{d}\sigma/\mathrm{d}t \propto e^{Bt}$, with diffractive slopes $B = 4$, 3, and 5 GeV^{-2} for πN, KN, and NN, respectively. These slopes are *representative* rather than fitted: they are single constants of the right order for each system rather than functions of energy fitted channel by channel to angular-distribution data. Their provenance is the standard compilations of the forward-elastic slope parameter B obtained from fits of $\ln(\mathrm{d}\sigma/\mathrm{d}t)$ to $A + Bt$: Ref. [37] for $\pi^{\pm}p$, $K^{\pm}p$, pp, and $\bar{p}p$ and Ref. [38] for the momentum dependence of B in $\pi^{\pm}p$ and $K^{\pm}p$ down to $p_{\mathrm{lab}} \simeq 0.3$ GeV/c. The constants adopted here reproduce the ordering $B_{KN} < B_{\pi N} < B_{NN}$ found in those fits and are deliberately taken at the low end of the compiled range: the slopes fall well below their asymptotic high-energy values at the sub-GeV momenta that dominate the cascade, where the forward peak is weak. At the momenta relevant here, the peaking factor $\exp[2Bp^{*2}(\cos\theta^{*} - 1)]$ is close to unity, so the scattering is in any case nearly isotropic and the choice has little leverage; it redistributes the direction of an already-counted interaction and leaves the integrated survival probability unchanged. Absorption and production are treated as isotropic with an approximate sequential phase space. The model is designed for spectrum and efficiency studies rather than as a precision transport code.

A freshly produced meson is not yet a fully formed color singlet and cannot reinteract at the full hadronic cross section the moment it is created. The cascade, therefore, offers an optional formation zone: each produced meson free streams over a formation length $L_f = (p/m)c\tau_f$ before its interactions are enabled, with $p/m = \beta\gamma$ dilating the rest-frame formation time $c\tau_f$ to the laboratory frame, the prescription used in NuWro's cascade [12,32]. The formation zone is off by default; when enabled, the default $c\tau_f = 0.342$ fm follows the SKAT parametrization [39], and it suppresses the first interaction of the primary meson and of mesons made in inelastic production, leaving struck recoil nucleons untouched. Because it can recover only yield that the cascade was removing, its effect is strongly channel dependent; we scan it in Sec. III B and use that scan as a proxy for the spread between FSI models.

The cross sections the cascade uses are shown in Fig. 6. For pions and nucleons, we adopt the Metropolis tables [33,34] as used in NuWro. The strange and η systems have no such tabulation, so we build short interpolation tables from the measured totals in the same spirit. The $S = +1$ kaons (K^{+} and K^{0}) are nearly transparent in nuclear matter, with total cross sections of only ∼10–18 mb and charge exchange their sole inelasticity; the $S = -1$ antikaons interact strongly (∼30–100 mb, rising toward threshold) and are dominated by absorption through $\bar{K}N \to Y\pi$ [40,41]. The ηN cross section is modeled by the $N^{*}(1535)$ Breit-Wigner that sits directly above the production threshold, with $\eta N \to \pi N$ as the leading inelastic channel.

A neutral kaon is propagated as the strangeness eigenstate in which it is produced, not as a K_S^0/K_L^0 or an incoherent $K^0 - \bar{K}^0$ mixture. The decays $p \to \ell^{+}K^0$ and $n \to \bar{\nu}K^0$ emit a $|K^0\rangle$, which carries $S = +1$ and, therefore, interacts with the same nearly transparent $S = +1$ cross sections as the K^{+}; a $\bar{K}^0$ appears only if charge exchange makes one. This is the physically correct choice inside the nucleus: the strong interaction conserves strangeness, and the kaon crosses the residual nucleus in $\sim 10^{-23}$ s, some 12 orders of magnitude faster than the K_S^0 lifetime, so the weak $K^0 - \bar{K}^0$ mixing that eventually turns the produced state into a K_S^0/K_L^0 superposition has no time to develop while the meson is still inside the medium. Mixing and decay are applied afterward, outside the nucleus, when the escaping neutral kaon is optionally decayed. The consequence is quantitative and not small: with the incoherent-mixture approximation, half of the neutral-kaon flux is exposed to $\bar{K}N \to Y\pi$ absorption and the predicted K^0 survival drops to ∼60%, whereas the eigenstate treatment leaves the K^0 as transparent as the K^{+} (Sec. III B).

LUNAR ships a single cascade, and it is worth being explicit about what that means for the uncertainty budget, since final-state interactions turn out to be the leading correction for most channels. Other approaches to the same

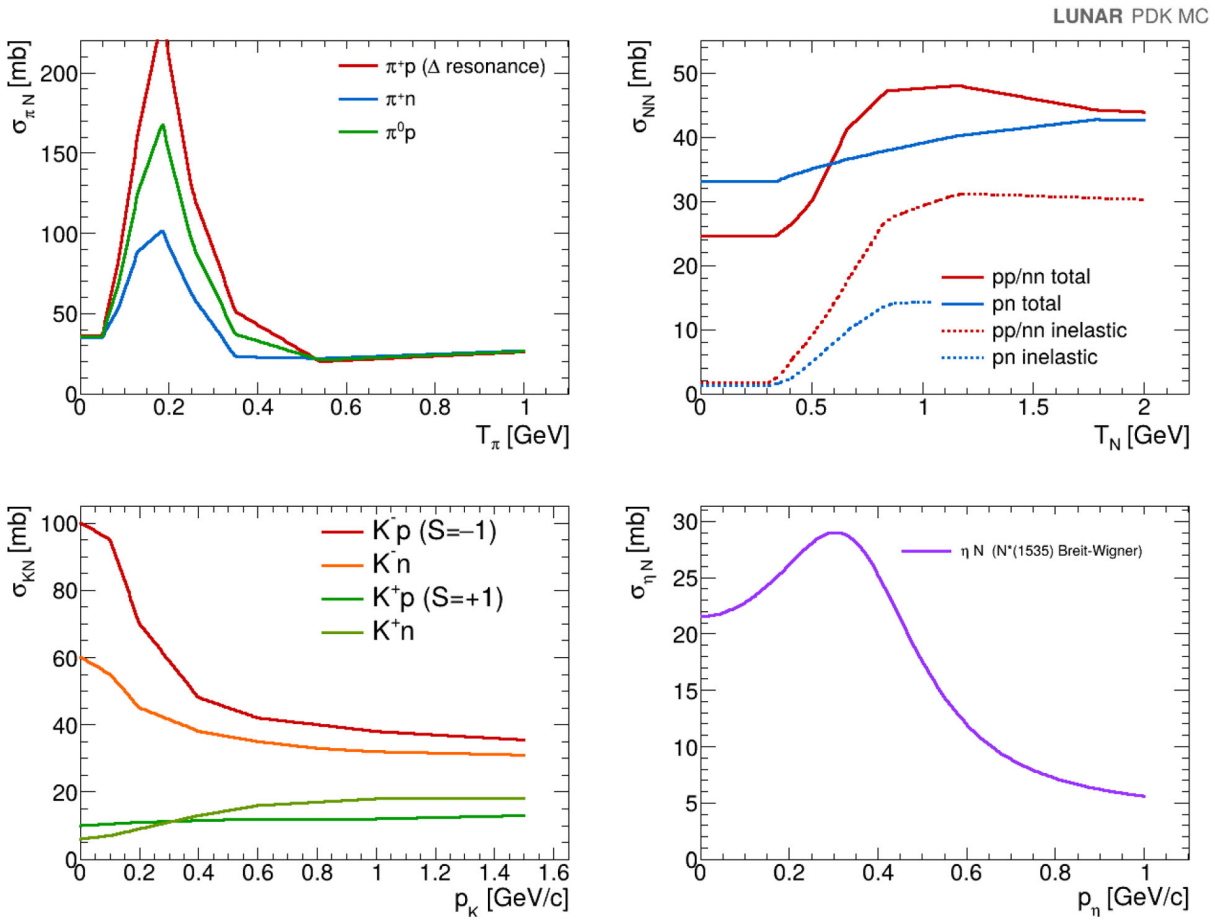


FIG. 6. Hadron-nucleon cross sections used by the LUNAR cascade: the $\Delta(1232)$ peak in πN (top left), the onset of NN inelasticity (top right), the nearly transparent K^+ contrasted with the strongly absorbed K^- (bottom left), and the $N^*(1535)$ resonance in ηN (bottom right).

physics differ substantially in construction: the intranuclear cascades of GENIE (hA and hN) [11] and NEUT trade a data-driven effective treatment against a microscopic one, the Liège cascade INCL++ [42] propagates all participants simultaneously with a mean-field potential, and GiBUU [13] replaces the cascade altogether by a Boltzmann-Uehling-Uhlenbeck transport with in-medium spectral functions and potentials, the approach used in the proton-decay study of Ref. [14]. A meaningful comparison would require running the same decays through those codes, which is beyond the scope of a dedicated generator paper. We instead bound the sensitivity from within LUNAR in two ways: by benchmarking the absorption probabilities against independent treatments (Table VI) and by scanning the formation zone (Table VII), which is the single ingredient that moves the survival probabilities most and which we carry as an explicit FSI-model band on the predicted event rates in Sec. III C. Interfacing LUNAR to an external cascade through a standard event record (Sec. V) would make the direct comparison straightforward, and we regard it as the natural next step.

III. RESULTS

A. Daughter spectrum: Fermi motion versus binding

The laboratory momentum of the decay meson is the experimental signature. A free proton at rest emits a kaon of fixed momentum $p_K^{\mathrm{free}} = (M_p^2 - m_K^2)/2M_p = 0.339\ \mathrm{GeV}/c$; in the nucleus, Fermi motion broadens this line and binding shifts it down through the reduced off-shell mass. The generator lets us vary one at a time and see which attribute of the spectrum each controls. All numbers in this subsection come from samples of 2×10^5 decays per configuration, so the statistical uncertainty on a mean is well below $1\ \mathrm{MeV}/c$ and every difference quoted is a model difference.

Varying the momentum model (Fig. 7) changes the spectral width. The mean-field models (the global and local Fermi gases, the harmonic-oscillator shell model, and the Gaussian) give the narrowest kaon peaks, with a root-mean-square spread of 63–67 MeV/c. The correlated and spectral-function models broaden it to 73–77 MeV/c, as their high-momentum tails feed a soft shoulder that pushes as much as ∼18% of kaons below $0.25\ \mathrm{GeV}/c$. The Bodek-Ritchie and correlated-Fermi-gas options, which carry the

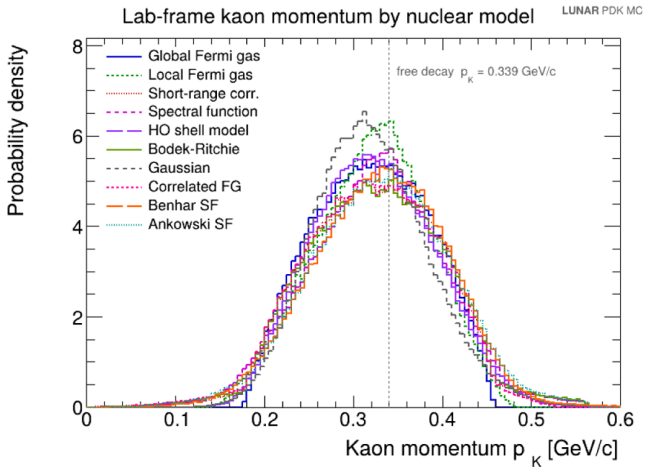


FIG. 7. Laboratory kaon-momentum spectrum for the ten physics momentum models (area-normalized), all with the default optical-potential binding; the dashed line marks the free-decay value $p_K = 0.339$ GeV/c. The pure Fermi-gas models give the narrowest peaks, while the correlated and spectral-function models broaden the spectrum and feed a low-momentum shoulder. The analytic ankowski spectral function is slightly broader than the tabulated benhar grid. Line styles as well as colors distinguish the models; models sharing a style differ in color.

most strength above k_F in Table II, produce the largest low-momentum tails, a direct mapping of initial-state high-momentum content onto the observable. The momentum model also moves the central value: the mean kaon momentum runs from 0.320 GeV/c for the Gaussian to 0.330 GeV/c for the two spectral functions, a spread of about 10 MeV/c.

Varying the binding model instead (Fig. 8 and Table IV) shifts the spectrum rigidly with almost no change in shape. The optical potential, being the deepest, yields the softest kaon; the constant and shell prescriptions lie close together and give slightly harder spectra. The displacement of the mean is ∼6 MeV/c, while the root mean square is fixed to better than 1 MeV/c.

TABLE IV. Laboratory kaon momentum versus binding model for two analytic momentum models (2×10^5 events each). "In-window" is the fraction of decays whose kaon momentum falls inside $[0.30, 0.38]$ GeV/c, i.e., within ± 0.04 GeV/c of the free-decay value, the containment proxy defined in Sec. III C. The binding choice moves the mean by ∼6 MeV/c and leaves the rms fixed, whereas changing the momentum model changes the rms by ∼20%. The spectral-function models are binding independent by construction.

Momentum model	Binding	$\langle p_K \rangle$ [GeV/c]	rms [GeV/c]	In-win. [%]
Local Fermi gas	Potential	0.321	0.063	45.4
	Constant	0.325	0.062	46.0
	Shell	0.327	0.063	46.0
Spectral fn. (toy)	Potential	0.325	0.075	41.2
	Constant	0.328	0.074	41.7
	Shell	0.330	0.074	41.9

The two nuclear ingredients are, therefore, *not* interchangeable in their effects, but neither are they orthogonal. The width of the kaon spectrum is controlled by the momentum model alone: the binding prescription leaves the rms unchanged, while going from a mean-field to a correlated or spectral-function momentum distribution widens it by ∼20%. The central value responds to both, and, importantly, it responds more strongly to the momentum model (∼10 MeV/c across the ten models) than to the binding prescription (∼6 MeV/c across the three). A meaningful nuclear uncertainty on the mean must, therefore, vary the momentum model; varying only the binding prescription,

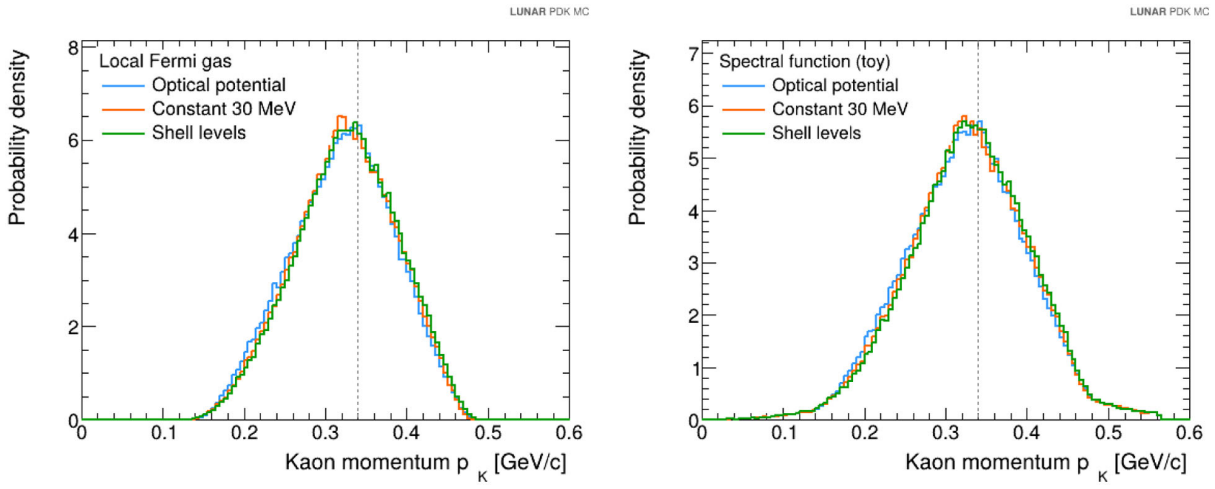


FIG. 8. Kaon-momentum spectrum across the three binding models, for the local Fermi gas (left) and the toy spectral function (right); the dashed line marks the free-decay value. The binding choice shifts the peak almost rigidly; the optical potential, being the deepest, gives the softest kaon and the broadest low-momentum shoulder.

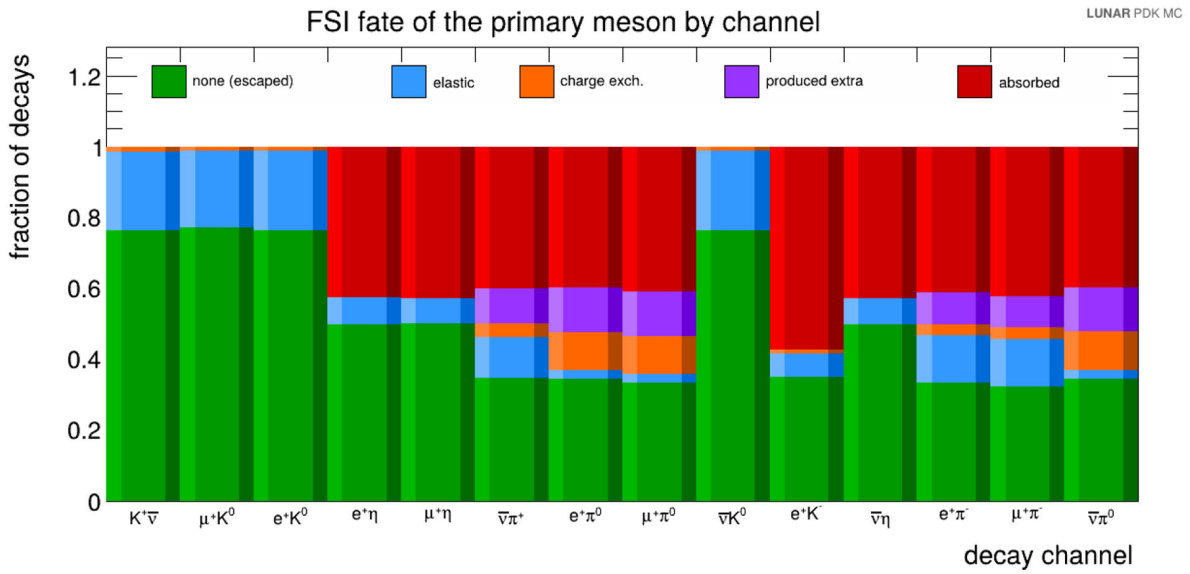


FIG. 9. Fate of the primary meson after the cascade in argon, for all 14 proton and neutron channels (each bar normalized to one decay). The $S = +1$ kaons are nearly transparent; the pions, the η, and the antikaon are strongly absorbed or charge exchanged. Local Fermi gas, 2×10^5 events for the six benchmark channels in Table V and 5×10^4 for the rest.

the more obvious knob, understates it by roughly a factor of 2. What can be said cleanly is the asymmetry: the binding model affects the central value and nothing else, whereas the momentum model affects everything.

B. Final-state interactions, channel by channel

The cascade reshapes the observable spectrum and, more importantly, removes signal mesons. We classify each event by the fate of the primary meson: it may escape untouched (none), scatter elastically, emerge as a different meson through charge exchange (CEX), survive while spawning additional mesons (produced), or fail to escape (absorbed). The first, second, and fourth of these leave a meson of the original species in the final state, and their sum is the quantity that propagates into every event rate we quote:

$$\varepsilon_{\mathrm{FSI}} \equiv P(\mathrm{none}) + P(\mathrm{elastic}) + P(\mathrm{produced}), \qquad (6)$$

the probability that the decay meson escapes the argon nucleus as itself. It is a survival probability, not a detector efficiency; charge exchange counts against it because the escaping meson is then a different particle from the one the search is looking for.

Figure 9 and Table V give these fractions for six benchmark channels, and the contrast between them is the central result of this section. The $S = +1$ kaons are almost transparent: three-quarters escape without interacting, the rest scatter elastically, and none is absorbed, since a K^+ or K^0 has no strangeness-conserving absorption channel on a nucleon. The pions are heavily reworked, with roughly 40% absorbed and a further quarter charge exchanged or augmented by produced mesons. The η and the antikaon are the most strongly absorbed (43% and 57%), the latter because its cross section grows toward threshold through $\bar{K}N \to Y\pi$.

TABLE V. Primary-meson outcome fractions, escaping signal-meson yield, and the shift of the mean observable momentum (cascade off → on, the latter averaged over escaping mesons of the signal species) for six benchmark channels in argon (2×10^5 events each, local Fermi gas). The yield is the survival probability $\varepsilon_{\mathrm{FSI}}$ of Eq. (6), used in the event-rate predictions in Sec. III C. Both $S = +1$ kaons are transparent: with strangeness conserved by the strong interaction, neither the K^+ nor the K^0 has an absorption channel on a nucleon.

Channel	Meson	None	Elastic	CEX	Prod.	Abs.	Yield	$\langle p \rangle$ shift [GeV/c]
$p \to K^+\bar{\nu}$	K^+	0.76	0.22	0.01	0.00	0.00	99%	0.321 → 0.295
$p \to \mu^+K^0$	K^0	0.77	0.22	0.01	0.00	0.00	99%	0.307 → 0.282
$p \to e^+\eta$	η	0.50	0.08	0.00	0.00	0.43	57%	0.293 → 0.279
$p \to e^+\pi^0$	π^0	0.34	0.02	0.11	0.12	0.40	49%	0.442 → 0.309
$p \to \bar{\nu}\pi^+$	π^+	0.35	0.12	0.04	0.10	0.40	56%	0.441 → 0.365
$n \to e^+K^-$	K^-	0.35	0.07	0.01	0.00	0.57	42%	0.321 → 0.311

TABLE VI. Benchmark of the primary-meson absorption probability in argon (from Table V) against independent treatments. Because the $\pi N/NN$ cross sections are the NuWro Metropolis tables and the K/η totals follow measured data, the comparison tests the cascade rather than its inputs.

Meson	S	LUNAR abs.	Independent treatment
π^0	0	0.40	∼0.40 (^{16}O, GiBUU) [14]
K^+	+1	0.00	Nearly transparent [15,40]
K^0	+1	0.00	Nearly transparent [40]
K^-	−1	0.57	Strongly absorbed [15,41]

The survival and absorption fractions in Table V rest on validated microphysics rather than free parameters: the πN and NN cross sections are the Metropolis tables [33,34] used by NuWro [12], embedded verbatim, and the K- and η-nucleon totals reproduce the measured data [40,41]. Table VI benchmarks the resulting meson absorption in argon against independent treatments: the π^0 absorption matches the GiBUU transport result of Yan *et al.* [14], and the $S = +1$ kaons are unabsorbed, lying below the $\bar{K}N \to Y\pi$ threshold that makes the $S = -1$ K^- opaque, as in the kaon survey of Barbu, Parvu, and Lazanu [15]. The only ingredients that are not taken from data are the elastic and charge-exchange angular slopes, which are representative constants rather than channel-by-channel fits (Sec. II E); because they redistribute only the direction of an interaction that has already been counted, the integrated survival $\varepsilon_{\mathrm{FSI}}$ is insensitive to them.

Table VII scans the surviving-meson yield against the formation time $c\tau_f$ of the formation zone described in Sec. II E. Because the formation zone can recover only yield that the cascade was removing, its impact is sharply channel dependent. The $S = +1$ kaon modes are unchanged (99% → 99%): the kaon is slow ($\beta\gamma \simeq 0.7$, $L_f \lesssim 0.3$ fm at the default) and already transparent, so there is nothing to recover. The strongly interacting pions shift the most (the $e^+\pi^0$ yield climbs from 49% with no formation zone to 60% at the default and to 81% at 1.0 fm and the $\bar{\nu}\pi^+$ from 56% to 66% and 83%) because the energetic pions ($\beta\gamma \sim 3$) free stream a sizable fraction of the nuclear radius before they can be absorbed; the slower η and $\bar{K}$ modes gain two to three points at the default and up to ∼7 percentage points at 1.0 fm. The formation zone is, thus, a leading systematic for the pion, η, and antikaon channels and a negligible one for the kaon searches. The predictions in Sec. III C use the default cascade (formation zone off), which is the conservative end of this range, and we carry the 0.342–1.0 fm spread in Table VII as an explicit FSI-model band on $\varepsilon_{\mathrm{FSI}}$ and, hence, on the predicted event rates.

TABLE VII. Formation-zone (FZ) sensitivity of the escaping signal-meson yield (the survival probability none + elastic + produced, i.e., the $\varepsilon_{\mathrm{FSI}}$ used in Sec. III C) for the six benchmark channels, as the formation time $c\tau_f$ is scanned. "no FZ" is the default cascade (the yield column in Table V); the formation zone monotonically raises the yield, by $\lesssim$1 point for the transparent $S = +1$ kaons but by tens of points for the strongly absorbed pions. The scan is run for all 14 channels and supplies the FSI-model band of Sec. III C; the six benchmarks are shown here. The same sample as Table V (2×10^5 events, local Fermi gas).

			$c\tau_f$ [fm]		
Channel	Meson	No FZ	0.342	0.5	1.0
$p \to K^+\bar{\nu}$	K^+	99%	99%	99%	99%
$p \to \mu^+K^0$	K^0	99%	99%	99%	99%
$p \to e^+\eta$	η	57%	59%	60%	63%
$p \to e^+\pi^0$	π^0	49%	60%	66%	81%
$p \to \bar{\nu}\pi^+$	π^+	56%	66%	70%	83%
$n \to e^+K^-$	K^-	42%	44%	46%	49%

The effect on the spectrum is shown in Figs. 10 and 11, where every curve is normalized per generated decay so that the cascade-on histogram exposes both the depletion (a smaller area, from absorption and charge exchange) and the softening, a low-momentum tail left by quasielastic energy loss. For the signal mode the kaon stays ∼99% transparent and merely softens slightly, its mean falling from 0.321 to 0.295 GeV/c. The $p \to e^+\pi^0$ pion, by contrast, drops to a 49% yield and grows the pronounced low-momentum bump that is the hallmark of pion final-state interactions. The same behavior carries over to the π^+ and η channels, while the K^0 follows the K^+. Across all modes, the surviving-meson yield (the yield column in Table V, whose components are shown channel by channel in Fig. 9) ranges from ∼99% for the $S = +1$ kaons down to ∼42% for the strongly absorbed K^-: the cascade is a negligible correction for the kaon searches but a leading effect for every other channel.

Because the cascade knocks nucleons loose and can add or charge exchange pions, the observable final state is richer than a single meson. Figure 12 shows the mean postcascade multiplicity of escaping protons, neutrons, pions, and kaons per decay. The K^+ mode leaves with about one kaon and almost no extra hadrons, while the pion and η modes eject two to four nucleons and a soft pion halo. This accompanying activity is directly relevant to event reconstruction in a liquid-argon detector, where additional short tracks and vertex activity around the decay candidate both aid identification and complicate it [43].

C. Expected event yields in DUNE

We now combine the nuclear-model and cascade outputs into expected event counts for DUNE. For simplicity, we assume two 10-kt liquid-argon modules (20 kt fiducial) and the technical-design exposure of 400 kt · yr, about 20 years [10]. Argon-40 supplies 2.71×10^{32} protons and 3.31×10^{32} neutrons per kilotonne, so a proton mode

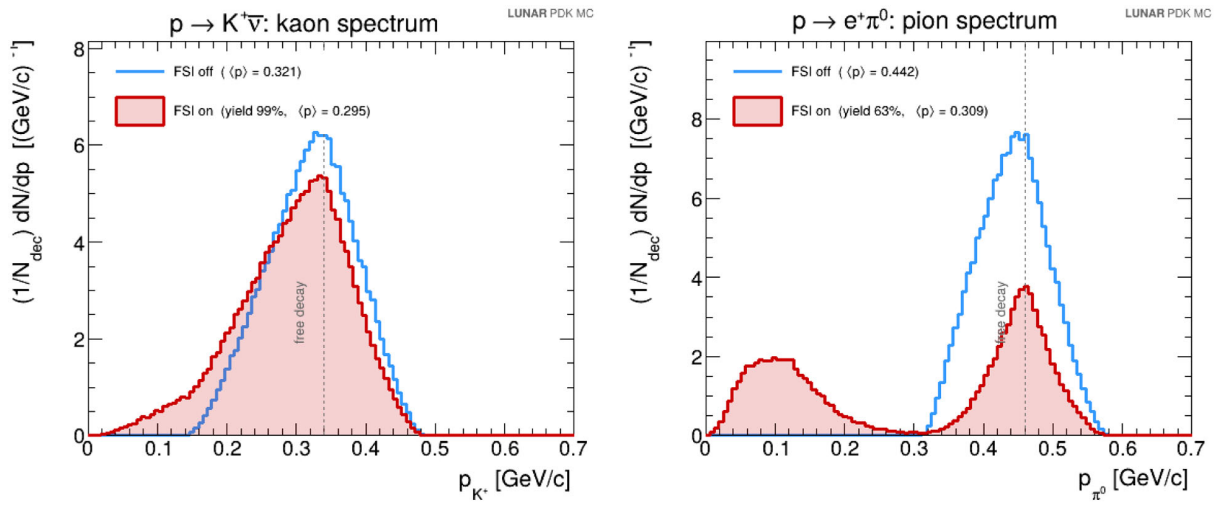


FIG. 10. Laboratory meson momentum with (red, filled) and without (blue) the cascade, normalized per generated decay, for the transparent K^+ (left) and the strongly interacting π^0 (right); the dashed line marks the free-decay value. The K^+ is barely affected, whereas the π^0 loses about half its yield and develops a large low-momentum bump. The legend quotes the integral of the "FSI on" histogram, which counts every meson of the signal species in the final state; for the π^0, this is 0.64 per decay rather than the primary-meson survival probability $\varepsilon_{\rm FSI} = 49\%$ in Table V, the difference being pions made in the cascade when the primary was absorbed or charge exchanged, which populate the low-momentum bump. For the transparent K^+, the two coincide.

draws on $N_p = 5.4 \times 10^{33}$ and a neutron mode on $N_n = 6.6 \times 10^{33}$ target nucleons. The expected count is

$$N = N_{\rm nuc}(1 - e^{-T/\tau})\varepsilon \simeq N_{\rm nuc}\frac{T}{\tau}\varepsilon \quad (\tau \gg T), \qquad (7)$$

with the detection efficiency factorized in the same way for every mode:

$$\varepsilon = \varepsilon_{\rm det}\varepsilon_{\rm FSI}, \qquad (8)$$

where $\varepsilon_{\rm FSI}$ is the cascade survival of Eq. (6), applied exactly once, and $\varepsilon_{\rm det}$ is a flat 30% reconstruction assumption, the same for all 14 channels. We deliberately do not substitute the mode-specific DUNE technical-design efficiencies quoted for selected channels [10]. Those DUNE numbers are evaluated on a simulation that *already* includes Fermi motion and final-state interactions, the $p \to e^+\pi^0$ value being limited by π^0 reabsorption, so they are not $\varepsilon_{\rm det}$ in the sense of Eq. (8) and cannot be multiplied by $\varepsilon_{\rm FSI}$ without double counting the cascade. A single flat $\varepsilon_{\rm det}$ keeps the comparison honest, leaves the golden kaon mode essentially where the DUNE number puts it ($0.30 \times 0.99 \simeq 0.30$), and is the more conservative choice for the π^0 modes. It also means the counts scale trivially: an analyst with a better per-channel reconstruction efficiency can rescale any row by $\varepsilon_{\rm det}/0.30$. For each mode, we set the lifetime τ to the current Super-Kamiokande 90% confidence

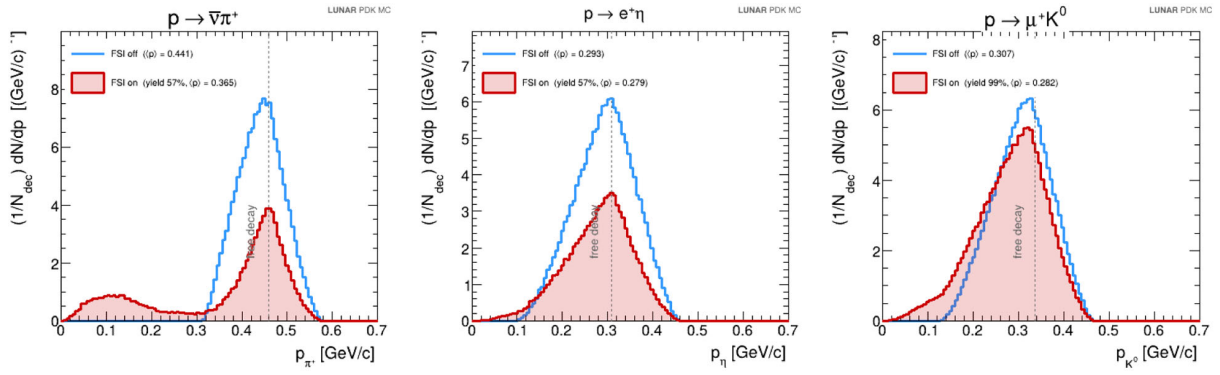


FIG. 11. Meson momentum with (red) and without (blue) the cascade for the π^+ (left), η (middle), and K^0 (right) channels, normalized per generated decay. The π^+ is depleted and softened most strongly and the η substantially; the $S = +1$ K^0 keeps essentially all of its yield and only softens, tracking the K^+ in Fig. 10.

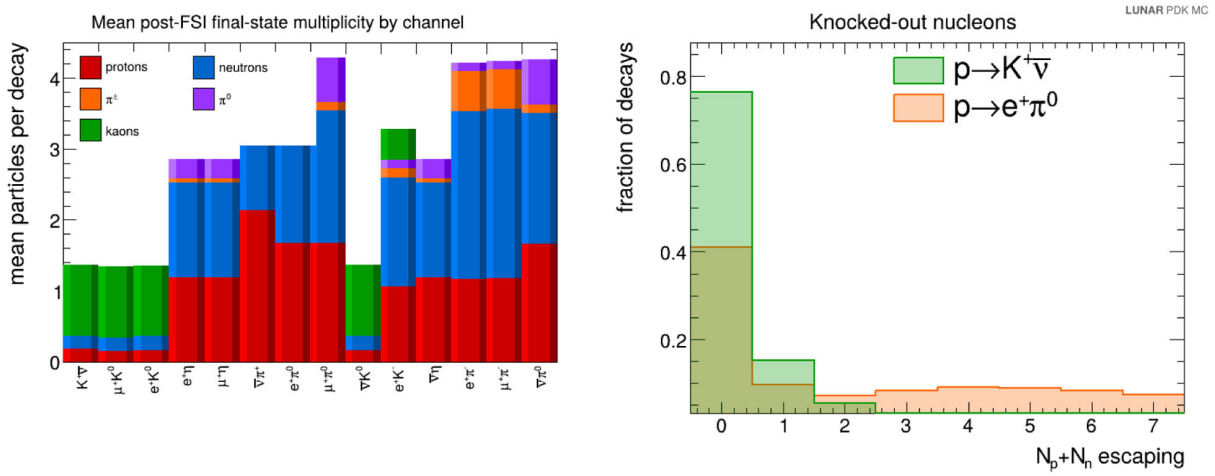


FIG. 12. Left: mean number of escaping protons, neutrons, charged and neutral pions, and kaons per decay, by channel. Right: the knocked-out nucleon multiplicity for the transparent $p \to K^+\bar{\nu}$ and the pion-cascade-dominated $p \to e^+\pi^0$.

level (CL) lower limit on τ/B [44], so that N is the largest yield consistent with present data, a discovery-reach ceiling that DUNE would fall below if the true lifetime is longer.

Anchoring $\varepsilon_{\rm det}$ to the DUNE assumption this way rests on a cascade consistent with DUNE's own: our K^+ survival (∼99%, Table V) agrees with the ∼97% that the GENIE-based DUNE simulation finds for the kaon escaping the argon nucleus, and our π^0 result, roughly half the pions reworked, is consistent with the DUNE $p \to e^+\pi^0$ efficiency being limited by inelastic intranuclear scattering [10]. The cascade survival $\varepsilon_{\rm FSI}$ runs from 0.99 for the $S = +1$ kaons to ∼0.42 for the K^- and is the dominant channel-to-channel modulation of the predicted counts.

Two uncertainty bands accompany each count, and it is worth being explicit about what they represent, because neither is a complete systematic. The first is the *nuclear-model* band. A selection in a liquid-argon detector will not be a cut on momentum; it will use $\mathrm{d}E/\mathrm{d}x$, track topology and the decay chain of the stopping meson. Any such selection is, nonetheless, tuned on a simulated signal spectrum, and it is the model dependence of that spectrum we want to expose. We, therefore, use a deliberately simple, transparent stand-in: the *signal-window containment efficiency* $\varepsilon_{\rm win}$, defined as the fraction of decays whose hadron-daughter laboratory momentum falls within ± 0.04 GeV/c of the free-decay value $p_{\rm free}$ for that channel. It is not an absolute reconstruction efficiency and is not applied as one; only its *relative* spread across the ten momentum models is used, as a proxy for how much a spectrum-based selection efficiency would move if the nuclear model were changed. Figure 13 shows $\varepsilon_{\rm win}$ channel by channel; the spread is $^{+11}_{-8}$ % about the model average, and we carry it as a relative

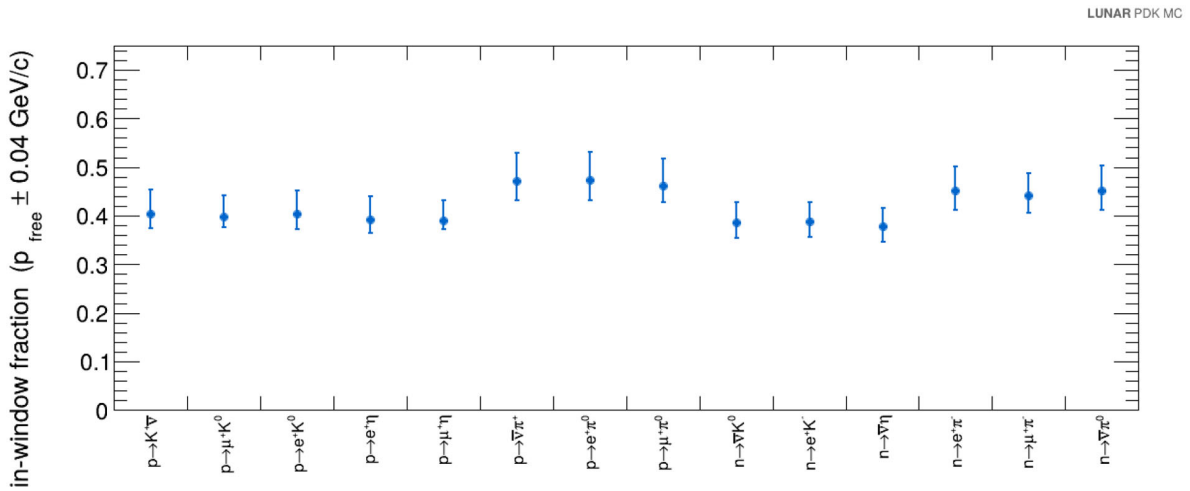


FIG. 13. Signal-window containment efficiency $\varepsilon_{\rm win}$, defined in the text, by channel. The marker is the mean over the ten physics momentum models and the bar their full spread. Only the relative spread is used, as a proxy for the nuclear-model dependence of a spectrum-based selection, and it is what sets the nuclear-model band on the event rates in Table VIII. Each of the 14 channels is generated with each of the ten momentum models, 5×10^4 events per combination, cascade off.

TABLE VIII. Predicted nucleon-decay events in DUNE (20 kt, 400 kt · yr) with τ set to each mode's current Super-Kamiokande 90% CL limit, so the counts are upper bounds allowed by present data. Every mode uses the same factorization $\varepsilon = \varepsilon_{\rm det}\varepsilon_{\rm FSI}$ with a flat $\varepsilon_{\rm det} = 0.30$. The quoted uncertainty on the FSI-on count is the nuclear-model band in Fig. 13. "FZ max" is the upper edge of the FSI-model band, obtained with the formation zone at $c\tau_f = 1.0$ fm (Table VII); it exceeds the nuclear-model band for every pion mode. The last column is the count with the cascade off ($\varepsilon_{\rm FSI} = 1$), so the gap to the FSI-on column is what the cascade removes. $n \to e^+K^-$ has no dedicated limit. Limit values follow the PDG compilation [44].

Mode	N	τ/B limit [yr] (90% CL)	$\varepsilon_{\rm det}$	$\varepsilon_{\rm FSI}$	$N_{\rm evt}$ (FSI on)	$N_{\rm evt}$ (FZ max)	$N_{\rm evt}$ (FSI off)
$p \to K^+\bar{\nu}$	p	5.9×10^{33} [26]	0.30	0.99	$5.4^{+0.7}_{-0.4}$	5.4	5.5
$p \to \mu^+K^0$	p	3.6×10^{33} [27]	0.30	0.99	$8.9^{+1.0}_{-0.5}$	8.9	9.0
$p \to e^+K^0$	p	1.0×10^{33} [25]	0.30	0.99	32^{+4}_{-2}	32	33
$p \to e^+\eta$	p	1.4×10^{34} [31]	0.30	0.57	$1.3^{+0.2}_{-0.1}$	1.5	2.3
$p \to \mu^+\eta$	p	7.3×10^{33} [31]	0.30	0.57	$2.5^{+0.3}_{-0.1}$	2.8	4.5
$p \to \bar{\nu}\pi^+$	p	3.9×10^{32} [29]	0.30	0.56	47^{+6}_{-4}	69	83
$p \to e^+\pi^0$	p	2.4×10^{34} [28]	0.30	0.49	$0.67^{+0.08}_{-0.06}$	1.1	1.4
$p \to \mu^+\pi^0$	p	1.6×10^{34} [28]	0.30	0.48	$0.98^{+0.12}_{-0.07}$	1.6	2.0
$n \to \bar{\nu}K^0$	n	1.3×10^{32} [25]	0.30	0.99	302^{+31}_{-25}	302	306
$n \to \bar{\nu}\eta$	n	1.6×10^{32} [30]	0.30	0.57	142^{+15}_{-12}	155	248
$n \to \bar{\nu}\pi^0$	n	1.1×10^{33} [29]	0.30	0.49	18^{+2}_{-2}	29	36
$n \to \mu^+\pi^-$	n	3.5×10^{33} [30]	0.30	0.55	$6.2^{+0.6}_{-0.5}$	9.3	11
$n \to e^+\pi^-$	n	5.3×10^{33} [30]	0.30	0.56	$4.2^{+0.5}_{-0.4}$	6.2	7.5
$n \to e^+K^-$	n	...	0.30	0.42	...	...	...

band on ε and, hence, on N. Because the spread enters as a ratio, treating it as a band on $\varepsilon_{\rm det}$ or on N is the same statement; we quote it on N because that is the quantity of interest.

The second is the *FSI-model* band. Since LUNAR provides one cascade (Sec. II E), we bound its modeling uncertainty with the formation-zone scan in Table VII, which is the ingredient that moves $\varepsilon_{\rm FSI}$ most. The predictions use the default cascade (formation zone off) as the central value, and the band extends upward to the $c\tau_f = 1.0$ fm result. For the pion modes, this band is several times wider than the nuclear-model band, the point at which the FSI treatment, not the nuclear ground state, becomes the limiting systematic.

Neither band includes the binding prescription (which shifts the central value of the spectrum, not the rate), the value of $\varepsilon_{\rm det}$ itself, or the atmospheric-neutrino background. The counts N are, thus, model-band estimates of the detected signal under a fixed, uniform efficiency assumption, not a complete systematic budget.

The results are collected in Table VIII and Fig. 14. Even saturating the present limits, DUNE would gather only a handful of $p \to K^+\bar{\nu}$ events ($\sim$5 in 400 kt · yr), the golden mode, for which the cascade matters at the percent level. The channels with the weakest current limits, $n \to \bar{\nu}K^0$, $n \to \bar{\nu}\eta$, and $p \to \bar{\nu}\pi^+$, have the most room for discovery, with tens to a few hundred allowed events. For the pion and η modes, the decisive systematic is the cascade, which roughly halves the rate, an effect several times larger than the $^{+11}_{-8}$ % nuclear-model band, and whose own modeling uncertainty, bounded here by the formation-zone band, is itself larger than that band: for $p \to e^+\pi^0$ the count moves from 0.67 to 1.1 across the formation-zone range, against ± 0.07 from the nuclear model. Where the FSI treatment is uncertain at this level, refining the nuclear ground state buys little. The kaon channels are barely touched by either: with the $S = +1$ kaons transparent, both $p \to K^+\bar{\nu}$ and the K^0 modes are limited by the nuclear model rather than the cascade. Final-state interactions are, thus, simultaneously a leading correction and a strongly mode-dependent one, and they reinforce the strange channels, $p \to K^+\bar{\nu}$ above all, as the cleanest nucleon-decay probes available to a liquid-argon detector.

A recent GiBUU study by Yan *et al.* [14] reaches a complementary conclusion. Treating $p \to e^+\pi^0$ in water Cherenkov detectors with full Boltzmann transport, they find the pion final-state-interaction uncertainty to be moderate, while the choice of Fermi-momentum distribution dominates the systematic, specifically on the atmospheric-neutrino background rate, a quantity our signal-only generator does not address. The two pictures are consistent rather than contradictory: both identify the nuclear initial state and final-state interactions as the controlling nuclear effects, and which one dominates depends on the target (argon versus water), the channel, and whether the limiting quantity is the signal efficiency and spectrum or the background rate. Our argon, signal-side scan across all channels puts the FSI cascade foremost, since it is what separates the transparent K^+ from the heavily reworked pions, whereas their water, background-side study of the

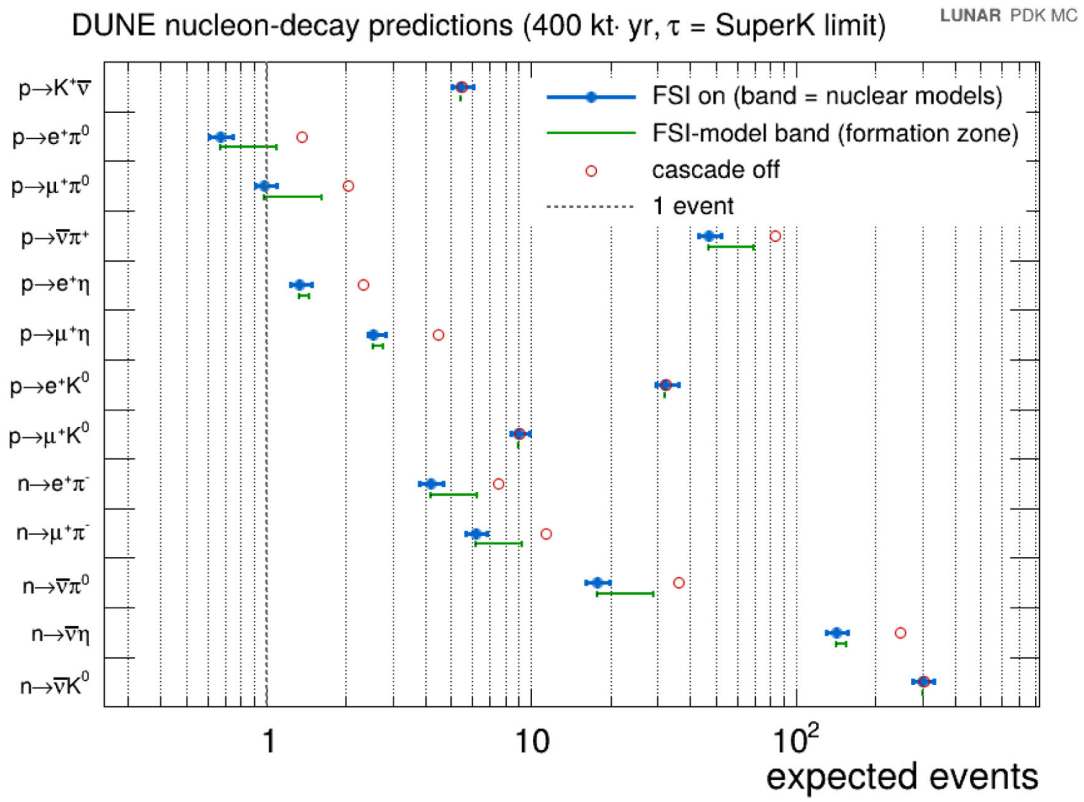


FIG. 14. Predicted events per channel in 400 kt · yr at τ equal to the current Super-Kamiokande limit. Filled markers carry the nuclear-model band (thick, blue); the thin green bar below each point is the FSI-model band, running from the default cascade up to the formation zone at $c\tau_f = 1.0$ fm. Open markers are the FSI-unfolded counts ($N/\varepsilon_{\mathrm{FSI}}$), so the gap to the filled marker is the cascade suppression: substantial for the pions and the η and negligible for the $S = +1$ kaons, whose two bands and two markers coincide. The dashed line marks one event.

single $e^+\pi^0$ mode foregrounds the Fermi-momentum distribution. Much of the contrast is one of metric: their quoted final-state-interaction uncertainty is the *variation* under FSI-model changes rather than its absolute size, and that absolute size (~40% of pions absorbed for production momenta of 0.4–0.5 GeV/c), in fact, agrees with our π^0 result. Barbu, Parvu, and Lazanu [15] survey the same Fermi-motion and FSI effects across argon, xenon, and water for $p \to \pi^+\bar{\nu}$ and $p \to K^+\bar{\nu}$ and reach the same qualitative ordering, the kaon far less affected than the pion, an independent cross-check of the cross sections and mean-free paths our cascade encodes; their analytic full-nucleus attenuation, with isospin-averaged K–N cross sections and a harder argon Fermi momentum (241 versus our 217 MeV/c), nonetheless, leaves the argon K^+ somewhat less transparent than our data-driven cascade.

IV. USING THE GENERATOR

LUNAR is intended to be picked up and used quickly. The build has no dependencies beyond a C++17 compiler, and the supplied pipeline compiles the generator, produces events, fills a ROOT tree, and draws the standard spectra in a single command. A minimal session selects a channel, a momentum model, a binding model, and the cascade:

```
./run.sh 100000 1 lfg potential pToKnu
```

generates 10^5 $p \to K^+\bar{\nu}$ events with the local Fermi gas and the optical potential. The momentum model (–model), binding model (–binding), channel (–channel), and cascade (–FSI on-off) are all chosen at run time; the complete option list and the output format are given in Appendix B. With the cascade enabled the generator writes a full postinteraction final state, one row per escaping particle, with its species, laboratory four-momentum, and a label recording the fate of the primary meson, so that downstream analyses see the same hadronic activity a detector would.

This lightness is the point: where the general-purpose generators carry sophisticated nuclear models but are built around the neutrino-interaction vertex and are heavy to deploy for a nucleon-decay study, LUNAR does one thing and exposes every ingredient. Three uses follow naturally. First, signal efficiency: because the generator reports the surviving meson and the accompanying nucleons and pions, it provides the per-channel acceptance and the hadronic activity that drive a liquid-argon selection.

Second, systematics: the ten momentum models and three binding prescriptions span the nuclear-model space, and the formation-zone scan brackets the cascade, so an analyzer can read off defensible uncertainty bands, as in the $^{+11}_{-8}$ % nuclear-model spread and the wider FSI-model band in Sec. III C, rather than guess them. Third, fast turnaround and pedagogy: a full model scan runs in seconds on a laptop, which makes the tool suitable both for rapid sensitivity estimates and for teaching the nuclear physics of nucleon decay.

V. POSSIBLE EXTENSIONS

The modular, header-only design makes LUNAR straightforward to extend, and several directions are natural. The nuclear input is the most localized: changing the density parameters, Fermi momenta, and shell tables retargets the generator to a different nucleus, so the same machinery could serve the oxygen of water-Cherenkov detectors, for direct cross-checks against Super-Kamiokande and Hyper-Kamiokande, or other liquid-argon variants, limited only by the availability of a tabulated spectral function.

The decay vertex can be broadened beyond the two-body modes treated here. Three-body and radiative final states, and the inclusive cascade decays of heavier mesons, would extend the channel list, while dinucleon decay and neutron-antineutron oscillation [9] would let the same nuclear and cascade machinery address baryon-number violation by two units. Each amounts to adding an entry to the channel registry and the corresponding phase-space generator.

The cascade itself admits systematic refinement. An in-medium K^+ potential, additional inelastic channels, and diffractive slopes fitted channel by channel would sharpen the strange-sector transport; a complete treatment of the secondary decays of escaping π^0, η, K^0_L, and hyperons, together with the deexcitation photons and nucleons emitted as the excited residual nucleus relaxes (extra activity that a liquid-argon detector can tag for background rejection and that must be modeled for realistic efficiency estimates) would complete the visible final state. On the nuclear side, fitted rather than representative proton and neutron Fermi momenta and shell energies would tighten the model band.

Finally, interoperability would broaden the generator's reach. Writing events in a standardized record such as HepMC3 [45] or the neutrino-specific NuHepMC [46] would let LUNAR feed detector simulations directly, and a thin interface to a liquid-argon simulation chain would close the loop between the nuclear physics modeled here and a full detector response. None of these requires structural change; each is a local addition enabled by the present design.

VI. SUMMARY

We have presented LUNAR, a compact and openly available Monte Carlo generator for bound-nucleon decay in argon-40. It samples the decaying nucleon from one of ten nuclear momentum models, binds it off the mass shell with one of three removal-energy prescriptions (including the momentum-dependent optical potential of Ref. [17]), performs the off-shell two-body decay into the laboratory frame, and transports the daughter meson out of the nucleus through a semiclassical intranuclear cascade, writing a complete final state for any of the standard proton and neutron channels.

Using the generator, we have separated the nuclear effects that shape the observable meson spectrum, finding that the momentum model alone controls its width while both the momentum and the binding model move its central value, the former by the larger amount; we have quantified final-state interactions mode by mode, showing that the nucleus is nearly transparent to the $S = +1$ kaons but absorbs roughly half of the pions, η mesons, and antikaons; and we have translated current Super-Kamiokande limits into expected DUNE yields across all standard channels, where the cascade emerges as the dominant systematic for every mode except the strange ones. These results reinforce $p \to K^+\bar{\nu}$ as the cleanest liquid-argon nucleon-decay probe and, for the other modes, identify final-state interactions, and the choice of FSI model rather than the nuclear ground state as the leading correction.

The generator has the expected limitations of a fast, dedicated tool. Most of the analytic momentum models are parametrized rather than fitted; the cascade uses tabulated cross sections and representative angular slopes rather than a channel-by-channel fit; and the detector response is not included. None of these is structural, and the modular design we have emphasized makes each a candidate for the refinements in Sec. V. We release LUNAR to the community in that spirit, as a transparent, extensible instrument for the nuclear physics of nucleon decay in liquid argon.

ACKNOWLEDGMENTS

This work was supported by the Science and Technology Facilities Council (STFC) Lancaster EPP Consolidated Grant 2025–2029: No. UKRI2846.

DATA AVAILABILITY

The results are based on software written by the author, i.e., LUNAR PDK Monte Carlo generator, which is publicly available [47]. All appropriate references are included as well.

APPENDIX A: MOMENTUM-MODEL PARAMETERS

The analytic momentum models in Table I use the following choices. The short-range-correlation (SRC) and toy spectral-function (SF) models place a default 20% of nucleons in a $1/p^4$ tail extending to $k_{\max} = 0.65$ GeV$/c$. The Bodek-Ritchie gas (BR) [18] uses a larger 25% fraction reaching to 1 GeV$/c$ and the correlated Fermi gas (CFG) a

20% contact tail to $2k_F$. The harmonic-oscillator shell model (HOSM) builds $n(p)$ from oscillator wave functions using the (n, l) quantum numbers of the argon shell table, and the Gaussian (GAUSS) is a smooth baseline with $\sigma = k_F/\sqrt{5}$ and no Fermi surface. The argon nuclear parameters [the density of Eq. (1), the Fermi momenta $k_F^p = 217$ MeV/c and $k_F^n = 230$ MeV/c, the isospin fractions, and the proton and neutron shell levels used by the shell binding model] follow Ref. [17]. The benhar option instead draws (p, E_{rem}) jointly from the tabulated argon spectral function of Refs. [19,20] and ankowski from an analytic effective spectral function (a $f_{corr} = 0.20$ correlated $1/p^4$ tail plus a shell mean field built from the argon orbitals); both ignore the binding parameters above.

APPENDIX B: COMMAND-LINE OPTIONS AND OUTPUT FORMAT

The generator is driven by the executable LunarPDKGenerator with the following options:—events (number of decays),—channel (mode key in Table III),—model (momentum model in Table I),—binding (potential, constant, or shell),—FSI (on or off),—decay-mesons (decay escaping π^0, η, and K^0),—formation-zone and—formation-time (free-stream-produced mesons one formation length before reinteraction),—seed, and—sf-file (override the spectral-function grid).

The—sf-file option accepts any tabulated $S(p, E)$ in the NuWro gridfun2d plain-text format, so the benhar model is not restricted to the argon grids shipped with the code. The file begins with three header lines (the number of removal-energy points and momentum points, the two axis minima, and the two axis maxima, all energies and momenta in MeV) followed by one block per momentum value: the momentum, then the (E, S) pairs at that momentum, energy index running fastest. Absolute normalization is irrelevant, as only ratios enter the sampling. Supplying a grid for a different nucleus is, therefore, enough to retarget the initial state, provided the corresponding density and Fermi-momentum parameters are updated as well (Sec. V).

With the cascade on, the output is one row per escaping particle (event, PDG, the laboratory momentum components px, py, and pz, the energy E, and a fate label outcome), preceded by a per-event header recording the nucleon momentum, removal energy, and primary-meson fate. With the cascade off, the generator writes the legacy momentum-only table (event, nucleon_p, d1_p, d2_p, e_rem) read by the single-distribution analysis macros. All momenta and energies are in GeV.

[1] H. Georgi and S. L. Glashow, Phys. Rev. Lett. **32**, 438 (1974).
[2] J. C. Pati and A. Salam, Phys. Rev. D **10**, 275 (1974); **11**, 703(E) (1975).
[3] S. Weinberg, Phys. Rev. Lett. **43**, 1566 (1979).
[4] F. Wilczek and A. Zee, Phys. Rev. Lett. **43**, 1571 (1979).
[5] N. Sakai and T. Yanagida, Nucl. Phys. **B197**, 533 (1982).
[6] S. Dimopoulos, S. Raby, and F. Wilczek, Phys. Lett. **112B**, 133 (1982).
[7] P. Nath and P. Fileviez Pérez, Phys. Rep. **441**, 191 (2007).
[8] K. S. Babu *et al.*, arXiv:1311.5285.
[9] D. G. Phillips II *et al.*, Phys. Rep. **612**, 1 (2016).
[10] B. Abi *et al.* (DUNE Collaboration), arXiv:2002.03005.
[11] C. Andreopoulos *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **614**, 87 (2010).
[12] H. Prasad *et al.* (NuWro Collaboration), arXiv:2501.11470.
[13] O. Buss *et al.*, Phys. Rep. **512**, 1 (2012).
[14] Q. Yan, A. Takenaka, K. Gallmeister, X.-G. Lu, U. Mosel, and Y. Zheng, Phys. Rev. D **113**, 095039 (2026).
[15] D. Barbu, M. Parvu, and I. Lazanu, Mod. Phys. Lett. A **40**, 2550038 (2025).
[16] S. Gardiner, Comput. Phys. Commun. **269**, 108123 (2021).
[17] C. Juszczak, J. A. Nowak, and J. T. Sobczyk, Eur. Phys. J. C **39**, 195 (2005).
[18] A. Bodek and J. L. Ritchie, Phys. Rev. D **24**, 1400 (1981); see also **23**, 1070 (1981).
[19] L. Jiang *et al.* (Jefferson Lab Hall A Collaboration), Phys. Rev. D **105**, 112002 (2022).
[20] R. D. Banerjee and A. M. Ankowski, Phys. Rev. D **109**, 073004 (2024).
[21] A. M. Ankowski and J. T. Sobczyk, Phys. Rev. C **77**, 044311 (2008).
[22] O. Benhar, A. Fabrocini, S. Fantoni, and I. Sick, Nucl. Phys. **A579**, 493 (1994).
[23] R. Subedi *et al.*, Science **320**, 1476 (2008).
[24] N. Fomin *et al.*, Phys. Rev. Lett. **108**, 092502 (2012).
[25] K. Kobayashi *et al.* (Super-Kamiokande Collaboration), Phys. Rev. D **72**, 052007 (2005).
[26] K. Abe *et al.* (Super-Kamiokande Collaboration), Phys. Rev. D **90**, 072005 (2014).
[27] A. Takenaka *et al.* (Super-Kamiokande Collaboration), Phys. Rev. D **106**, 072003 (2022).
[28] A. Takenaka *et al.* (Super-Kamiokande Collaboration), Phys. Rev. D **102**, 112011 (2020).
[29] K. Abe *et al.* (Super-Kamiokande Collaboration), Phys. Rev. Lett. **113**, 121802 (2014).
[30] K. Abe *et al.* (Super-Kamiokande Collaboration), Phys. Rev. D **96**, 012003 (2017).
[31] N. Taniuchi *et al.* (Super-Kamiokande Collaboration), Phys. Rev. D **110**, 112011 (2024).
[32] T. Golan, C. Juszczak, and J. T. Sobczyk, Phys. Rev. C **86**, 015505 (2012).

[33] N. Metropolis, R. Bivins, M. Storm, Anthony Turkevich, J. M. Miller, and G. Friedlander, Phys. Rev. **110**, 185 (1958).
[34] N. Metropolis, R. Bivins, M. Storm, J. M. Miller, G. Friedlander, and Anthony Turkevich, Phys. Rev. **110**, 204 (1958).
[35] L. L. Salcedo, E. Oset, M. J. Vicente-Vacas, and C. Garcia-Recio, Nucl. Phys. **A484**, 557 (1988).
[36] E. Oset and L. L. Salcedo, Nucl. Phys. **A468**, 631 (1987).
[37] T. Lasinski, R. Levi Setti, B. Schwarzschild, and P. Ukleja, Nucl. Phys. **B37**, 1 (1972).
[38] T. Lasinski, R. Levi Setti, and E. Predazzi, Phys. Rev. **179**, 1426 (1969).
[39] D. S. Baranov *et al.*, Rep. No. PHE 84-04, 1984.
[40] C. B. Dover and G. E. Walker, Phys. Rep. **89**, 1 (1982).
[41] E. Friedman and A. Gal, Phys. Rep. **452**, 89 (2007).
[42] D. Mancusi, A. Boudard, J. Cugnon, J.-C. David, P. Kaitaniemi, and S. Leray, Phys. Rev. C **90**, 054602 (2014).
[43] A. Bueno, A. J Melgarejo, S. Navas, Z. Dai, Y. Ge, M. Laffranchi, A. Meregaglia, and A. Rubbia, J. High Energy Phys. 04 (2007) 041.
[44] S. Navas *et al.* (Particle Data Group), Phys. Rev. D **110**, 030001 (2024).
[45] A. Buckley, P. Ilten, D. Konstantinov, L. Lönnblad, J. Monk, W. Pokorski, T. Przedzinski, and A. Verbytskyi, Comput. Phys. Commun. **260**, 107310 (2021).
[46] S. Gardiner, J. Isaacson, and L. Pickering (NuHepMC Collaboration), SciPost Phys. Codebases **57**, 1 (2025).
[47] https://github.com/LUNARPDK/lunarpdkē.