



Optimizing the SHiP neutrino detector using Key4hep

Search for Hidden Particles

Jesús P. Márquez Hernández






Introduction

Frontiers of Particle Physics

- ▶ Higgs discovery made the Standard Model “complete” (circa 2012)
- ▶ The SM is a great theory but does not represent the full picture
- ▶ New Physics should exist but we have no definitive predictions on the masses and coupling constants of new particles
- ▷ Yesterday we had some nice introductions (DM, [axions](#))




Properties of Hidden Particles/Hidden Sector

- ▶ HS production and decay rates are strongly suppressed wrt the SM
 - ▷ BR of $O(10^{-10})$
 - ▷ Long-lived objects
 - ▷ Very weak interaction with matter
 - ▷ May decay to very diverse final states

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \mathcal{L}_{\text{mediator}} + \mathcal{L}_{\text{HS}}$$

Portal models	Final states
HNL	$l^+\pi^-, l^+K^-, l^+p^-$
Vector, scalar, axion portals	l^+l^-
HNL	$l^+l^-\nu$
Axion portal	$\gamma\gamma$




See back-up for a more detail introduction: dark portals, HNLs, etc.

Main experimental challenge is **background suppression**:
 Full reconstruction + PID are essential



The Hidden Sector and Dark Matter

- Portals to a HS could explain not only DM, but also ν masses, the baryon asymmetry, or the strong CP problem





In green: within SHiP reach (SND scattering · HSdS decays)



The SHiP experiment

- ▶ Search for **Hidden Particles** & **neutrino physics**!
- ▶ Two main detector systems:
 - ▷ Scattering neutrino **D**etector (SD)
 - ▷ **H**idden **S**ector **D**ecay **S**pectrometer (HSDS)



Our work goes here,
optimizing the SD detector



Physics reach of SHiP

► Split in the two main detector categories:

▷ Hidden Sector Decay Spectrometre (HSDS)

▷ Scattering neutrino Detector (SD)

	Physics model	Final state
HSDS	SUSY neutralino	$\ell^\pm \pi^\mp, \ell^\pm K^\mp, \ell^\pm \rho^\mp, \ell^+ \ell^- \nu$
	Dark photons	$\ell^+ \ell^-, 2\pi, 3\pi, 4\pi, KK, q\bar{q}, D\bar{D}$
	Dark scalars	$\ell\ell, \pi\pi, KK, q\bar{q}, D\bar{D}, GG$
	ALP (fermion coupling)	$\ell^+ \ell^-, 3\pi, \eta\pi\pi, q\bar{q}$
	ALP (gluon coupling)	$\pi\pi\gamma, 3\pi, \eta\pi\pi, \gamma\gamma$
	HNL	$\ell^+ \ell'^- \nu, \pi l, \rho l, \pi^0 \nu, q\bar{q}' l$
	Axino	$\ell^+ \ell^- \nu$
	ALP (photon coupling)	$\gamma\gamma$
	SUSY sgoldstino	$\gamma\gamma, \ell^+ \ell^-, 2\pi, 2K$
SD	LDM	electron, proton, hadronic shower
	$\nu_\tau, \bar{\nu}_\tau$ measurements	$\tau^\pm$
	Neutrino-induced charm production (ν_e, ν_μ, ν_τ)	$D_s^\pm, D^\pm, D^0, \bar{D}^0, \Lambda_c^+, \bar{\Lambda}_c^-$



Neutrino Physics at SHiP

► The high neutrino flux allows for the measurement of strange quark nucleon content via charm production in ν interactions

	CC DIS (W HGC)	Charm CC DIS (W HGC)	CC DIS (Fe MGTC)	Charm CC DIS (Fe MGTC)
N_{ν_e}	6.8×10^4	4.1×10^3	1.6×10^5	9.8×10^3
N_{ν_μ}	2.0×10^5	8.7×10^3	4.6×10^5	2.0×10^4
N_{ν_τ}	2.1×10^3	1.2×10^2	5.1×10^3	2.8×10^2
$N_{\bar{\nu}_e}$	1.4×10^4	7.4×10^2	3.7×10^4	1.9×10^3
$N_{\bar{\nu}_\mu}$	4.4×10^4	1.7×10^3	1.1×10^5	4.1×10^3
$N_{\bar{\nu}_\tau}$	1.5×10^3	8.6×10^1	3.8×10^3	2.1×10^2

► A total of 26 ν_τ candidates have been observed, *ever*

► SHiP will collect $\sim 10k$ ν_τ over 15 years

• This allows for the determination of the never-before measured F_4 & F_5

$$\begin{aligned}
 \frac{d^2 \sigma^{\nu(\bar{\nu})}}{dx dy} = & \frac{G_F^2 M E_\nu}{\pi(1 + Q^2/M_W^2)^2} \left\{ (y^2 x + \frac{m_\tau^2 y}{2E_\nu M}) F_1 + \right. \\
 & \left[(1 - \frac{m_\tau^2}{4E_\nu^2}) - (1 + \frac{Mx}{2E_\nu}) y \right] F_2 \pm \left[xy(1 - \frac{y}{2}) - \frac{m_\tau^2 y}{4E_\nu M} \right] F_3 + \\
 & \left. \frac{m_\tau^2(m_\tau^2 + Q^2)}{4E_\nu^2 M^2 x} F_4 - \frac{m_\tau^2}{E_\nu M} F_5 \right\}.
 \end{aligned}$$



Neutrino observation at SHiP


► ν_μ : straightforward, nothing come in and a track comes out, tell apart parity using muon bending

► ν_e : more challenging, nothing come in and an electromagnetic shower takes place inside of your detector. Parity is complex to identify


► ν_τ : challenging as τ s have many possible decays, look for nothing coming in and a vertex associated to *multiple tracks*.

▷ Parity obtained by looking at charge of everything through the bending of tracks

complexity




SD baseline geometry



The SHiP timeline

- ▶ SHiP is an approved CERN experiment since 2024
- ▶ We (IJCLab, with the SiPad detector) joined about 1 year ago
- ▶ I was given the task to optimize the SD detector for the TDR
 - ▷ The software was in renewal process, far to be ready for an optimization study
 - I made my own SD software: “key4ship”

We are here





Software & optimization (work in progress)



Key4ship: main pipeline

- ▶ The full stack is *almost ready* to go for optimization
 - ▶ The reconstruction chain has:
 - ▷ Sequential ACTS tracking: Local (per-detector) and global (spliced) tracks
 - μ events: $\sim 95\%$ eff. ($\text{hits}_\mu \geq 20$), $\sim 70\%$ ($10 \geq \text{hits}_\mu \geq 20$), $\sim 20\%$ ($\text{hits}_\mu \geq 10$)
 - ▷ Clustering – CyberPFA-inspire (2D + 3D) & CLUE3D for comparison
 - ▷ PFlow-inspired reconstruction – μ track remove from clusters
- } WIP
- ▶ A few things are still missing:
 - ▷ Digitization
 - ▷ Realistic FOMs (\$ estimators and physics goals)




https://github.com/marherje/key4ship_PoC



SD baseline geometry (in key4ship)


- ▶ This geometry has $120 + 22 + (14+14+14)$
 - ▷ Set to the maximum number of layers per detector in their space budget
 - ▷ MTC is fixed due to magnet design constraints
- ▶ Remarks on MTC configuration
 - ▷ The iron plates are extended to 60 cm along the y-axis
 - ▷ And complemented with steel along the x-axis
 - Basically: **magnetic** vs **non-magnetic**



Tracking with **ats**: Example of local + spliced tracks

Spliced in all detectors

SiTarget local tracks



Optimization with MOBO


► What's Multi Objective Bayesian Optimization?

- Find the best trade-offs between multiple objectives *using less simulations*
- Surrogate model → learns from previously evaluated configurations
- Acquisition function → selects the most promising next configurations
- Iterative process → balances exploration & exploitation
- Pareto front → identifies the non-dominated configurations

► Why use it for detector layout optimization?


- Large parameter space → many possible detector configurations
- Expensive evaluations → each configuration requires a full simulation + reconstruction
- Multiple competing objectives → performance, cost, resolution, material availability, etc.
- No single “best” configuration → MOBO finds the Pareto-optimal trade-offs

► TLDR: 2 D (S vs B) → N D (Objectives) , ROC → Pareto front , AUC → Hyper-volume



Optimization with MOBO (Dummy test)

► Flow chart:



► Search space (8 D): SiTarget layers 20–120, SiPad layers 4–22, W thickness 2–8 mm (SiTarget) and 5–30 mm (SiPad), SiTarget layer slack 1–15 mm, SiPad layer gap 0.5–10 mm, SiTarget XY-plane gap fraction 0–1, SiTarget–SiPad separation 0.002–200 mm. Fixed: frame gap, module offset, SiPad envelope slack

► Conflicting objectives:

- ▷ Maximize fraction of events with ≥ 1 global track – Free exploration
- ▷ Maximize SiPad hits – Known trend (sanity check)
- ▷ **Minimize cost** (based on Si and W weight) – Known trend (sanity check)



ATTENTION



This is a dummy test to check the
 software implementation from
 simulation to optimization.
 The scores have no real meaning
 besides acting as a guide for us
 to check if MOBO is working

► Dummy test: 10k Genie events per trial, 200 trials (baseline + 18 Sobol + 181 qLogNEHVI)


- ▷ ~80h of wall time (11% Genie, 60% ddsim, 28% ACTS, 1% other) – *No clustering, no analysis*



Optimization with MOBO (Dummy test - results)


Phase	Completed	eff_track median (max)	nhits_sipad median (max)	cost proxy median (min)	Points on the Pareto front
Baseline	1	0.5465 (0.5465)	351 475 (351 475)	1585 (1585)	0
Sobol	18	0.5272 (0.5533)	205 925 (371 488)	924 (341)	3
qLogNEHVI	181	0.5405 (0.5688)	354 016 (392 440)	947 (410)	24
All	200	0.5387 (0.5688)	350 534 (392 440)	949 (341)	27

trial	track eff.	SiPad hits	cost proxy	ST layers	ST W [mm]	SP layers	SP W [mm]
t00340	0.5688	321 k	1471	85	8.00	20	24.8
t00270	0.5681	361 k	1176	64	8.00	21	19.2
t00251	0.5656	346 k	1136	57	8.00	21	22.9
t00214	0.5651	360 k	1090	44	8.00	22	30.0
t00184	0.5632	378 k	1125	47	8.00	22	30.0
t00191	0.5590	378 k	1009	43	8.00	22	24.5
t00268	0.5535	357 k	961	39	8.00	21	26.4
t00314	0.5523	356 k	868	32	8.00	20	27.6
t00237	0.5520	392 k	829	23	6.77	22	30.0
t00347	0.5381	388 k	706	28	8.00	22	14.5
t00030	0.5126	339 k	493	20	8.00	22	5.0
t00012	0.4744	49 k	341	25	3.69	4	17.7
baseline (t00000)	0.5465	351 k	1585	120	3.50	22	10.0




Dummy test example (Cheapest configuration)

- ▶ Missing segments due to lack of granularity



Dummy test example (Best configuration)

- Best configuration, best track splicing



Optimization with MOBO (Dummy test - conclusions)

- ▶ The Pareto front beats the baseline on **all three objectives at once**
 - ▷ And **sanity checks passed, MOBO works**
- ▶ The design direction is clear: fewer, thicker layers. The front lives at 20–85 SiTarget layers against the baseline's 120, with tungsten at maximum
 - ▷ Track efficiency comes mostly from muons reaching the MTC, not from more measurement points inside the SiTarget & SiPad helps cleaning ghost hits thus giving cleaner tracks
- ▶ The search space was badly bounded:
 - ▷ SiTarget_WThickness saturates at its 8 mm ceiling in 20 of the 27 front points
 - ▷ A lesson for future runs: One short pre-run for estimate the needed hyper-space
- ▶ The free exploration FOM was poorly designed - more in the next slide
 - ▷ Another lesson: MOBO can exploit weakness in your FOMs



A note on Functions Of Merits (FOMs)

- ▶ It's clear from these first tests that good FOMs will determine how good MOBO optimizes in the future, **they need to be well designed**
 - ▷ We also need to measure scoring uncertainties so better means *better for real* (TBD)
- ▶ I used “Fraction of events with $N_{\text{Tracks}} \geq 1$ ” as one of the scores, which turned out to be a *very weak indicator of tracking quality*
 - ▷ Moving from the best configuration to the cheapest yields a 16.5% lost of “tracking efficiency” while saving ~70% of the cost
 - In reality, the lost of track quality is overlooked by this scoring system
 - 5x fewer spliced tracks, longer track carries less $\frac{1}{2}$ the measurement points!

Variant	SiTarget	SiPad	Events with a spliced track	Longest track
t00270 — best choice	64 × 8.0 mm	21 × 19.2 mm	349	185 hits
t00237 — max SiPad hits	23 × 6.8 mm	22 × 30.0 mm	319	80 hits
t00012 — cheapest	25 × 3.7 mm	4 × 17.7 mm	64	74 hits



Conclusions & next steps

- ▶ Key4ship is in a fairly advanced phase!
 - ▷ Tracking tracks, optimizer optimizes, clustering being implemented 😊
- ▶ The remaining next steps are well defined:
 - ▷ Digitization
 - SiPad: Mostly solved from recent TB analysis, will be ported soon
 - SiTarget/MTC: I will need help from other SHiP colleagues
 - ▷ High-level analysis
 - Trivial to adapt C++/ROOT analysis to read from EDM4hep (But might need digitization first)
 - ▷ Realistic FOMs
 - Hard to design, easy to implement → O(minutes)
- ▶ The next next steps (after SD optimization):
 - ▷ Knowledge-transfer & contribution to the SHiP's software development (Phlex-based)





Thanks for your attention!
Merci de votre attention!
Vielen Dank für Ihre Zeit!






Back-up

Dark Matter candidates



Dark Matter Candidates: Very little clue on mass scales

Too small mass
⇒ won't "fit"
in a galaxy!



Slide from
H. Murayama
(ESPPU 2020)


In Particle Physics the range can go from $O(10^{-24})$ (ULDM) to $O(10^{20})$ (Q-balls)



Broad discovery reach: DM portals

► SHiP will unlock several orders of magnitude of several different portals in the phase space of 1-10 GeV

BC1 (dark photons)




Multidecay and
other models:

[arXiv:2312.14868](https://arxiv.org/abs/2312.14868)

[arXiv:2503.01760](https://arxiv.org/abs/2503.01760)

[arXiv:2510.23696](https://arxiv.org/abs/2510.23696)

BC4 (scalars without hSS coupling)




Spin correlations

[arXiv:2606.19543](https://arxiv.org/abs/2606.19543)

Lesya Shchutka,
SPC September 2026

BC5 (scalars with hSS coupling fixed such that $\text{Br}(h \rightarrow SS) = 0.01$)



Post-discovery


[arXiv:2312.05163](https://arxiv.org/abs/2312.05163)

[arXiv:2609.19269](https://arxiv.org/abs/2609.19269)



Broad discovery reach: HNLs

Discovery case for HNLs: more than counting



SHiP is uniquely sensitive in the region, impacting cosmological observations

[Akita, Ovchinnikov, Syvolap](#)

• What is it?

- mass, lifetime, flavour-dependent couplings

• Probe Dirac vs Majorana nature

- decay angular distributions $\rightarrow$ LNC vs LNV

• One state or a quasi-degenerate pair?

- observe spatial HNL oscillations $\rightarrow$ measure $\Delta M \sim 10^{-6} \text{ eV}$ in favourable scenarios

[Tastet & Timiryasov](#)

• Does it explain fundamental puzzles?

- flavour-dependent couplings $\rightarrow$ test the seesaw origin of neutrino masses
- $M, \Delta M, U_\alpha^2 + \text{PMNS } \delta_{\text{CP}} \rightarrow$ test whether the HNLs can generate the baryon asymmetry

[Drewes et al](#)

[Caputo, Hernandez et al](#)



[De Vries, Drewes et al](#)

5



Location at CERN

New 400 GeV proton beam line branched off the splitter section of the SPS transfer line to the North




Proton yield and beam delivery

- Nominal beam intensity 4×10^{13} pot per spill
- Baseline scenario: annual yield of 4×10^{19} pot to the BDF, and 10^{19} pot to the other experiments in the North Area, while respecting HL-LHC requirements
- SHiP sensitivities assume 5×10^{20} pot in five years of nominal operation


Sensitivity to dark photons & DM

- ▶ SHiP sensitive to Dark Photons and DM Fermions up to 5 GeV
- ▶ SHiP will dominate the phase space of (kinematically allowed) dark photon masses


Limits in ESPPU Physics Briefing Book on A'



Contributions to Dark Photon Limits (E.Svensson)




Limits on χ by SHiP (M. Climescu at Queen Mary)




For completeness I: Sensitivity to HNL

- Among projects studied by ESPPU SHiP and DUNE are the only approved experiments

Future limits as in ESPPU Briefing Book




Zoom into SHiP (2504.06692v1)




For completeness II: Sensitivity to Light Scalars

- ▶ SHiP sensitive to light scalars up to 5 GeV
- ▶ Significant increase reach if both scalars in pair production can be fully reconstructed!

Future limits as in ESPPU Briefing Book




Zoom into SHiP (2504.06692v1)



For completeness III: Sensitivity to ALP

- ▶ Below 5 GeV SHiP would/will dominate the landscape for decades
- ▶ Could only be topped by LCF beam dump facility

Future limits as in ESPPU Briefing Book



Our detector

► Sampling calorimeter “SiPad”:

- AKA CALICE SiW-ECAL, LUXE ECAL-p/e, SiPixel detector, etc.
- Mostly based at IJCLab (and LLR), but contributors all around the world
- For a comprehensive description: See [Roman's](#) or [Vincent's](#) slides

► Not constrained to be an ECAL


- Small pad size ($5.52 \times 5.52 \text{ mm}^2$) → suitable for tracking

► Basic ASU size of $18 \times 18 \text{ cm}^2$

- 4 Si sensors → 4 chips each → 64 channels each (1024 ch/ASU)
- Up to 5 ASUs in a single readout
 - (2x2) $36 \times 36 \text{ cm}^2$, (3x3) $54 \times 54 \text{ cm}^2$, ..., layers can be mounted

► Good shower & track separation, and PID capabilities

- Already studied for LUXE ([10550/108985](#))



3x3 ASUs arrays



Genie events in key4ship

- ▶ A basic pipeline [geometry (gdml) → Genie events (root)] was provided by SHiP's Genie expert (A. Iuliano), I ported into key4ship
- ▶ Current Genie pipeline in key4ship:
 - ▷ Export DD4hep compact xml to gdml
 - ▷ Parse gdml to Genie stack
 - SPS flux, all six (anti-)neutrino flavors mixed
 - Tune and splines (sim models and cross sections per nucleus) are preconfigured
 - Generate ROOT file with gst ntuples, normalised in PoT
 - ▷ The ROOT ntuples are converted to HepMC with a python script
 - ▷ HepMC is parsed to ddsim
 - ▷ From there, standard key4ship's Gaudi pipeline



Tracking with : Sequential implementation

1. Measurements — SiTarget 1D strips (X/Y planes), SiPad 2D pads, MTC SciFi U/V $\pm 5^\circ$ stereo
2. Seed cleaning — Multiple ghost and shower hits filtering
3. Event topology — Scaling of min-hit requirements w.r.t. the beginning of the interaction
4. Seeding, independent per detector — Tracklet Hough transform
5. Prototracks — Combinatorial Kalman Filter (CKF)
6. Duplicate removal — best-first by χ^2/ndf , drop candidates sharing >70 % of hits or cellIDs
7. Intra-detector merging — collinear segments of the same detector are joined and refitted
8. Inter-detector splicing — straight extrapolation to the SiTarget|SiPad|MTC boundaries, position/direction window, combined refit as final judge $\rightarrow$ GlobalTracks
9. Promotion — a good single-detector segment with no neighbor to splice to (e.g. interaction inside the MTC) is promoted to a global track



Tracking with : Some preliminary numbers


Event class	n	≥ 1 global	tracks/ev	0	1	2	3	≥ 4
Muon events	1117	77.9 %	0.99	22.1 %	64.2 %	9.9 %	1.7 %	2.1 %
— muon ≥ 20 hits	733	95.0 %	1.21	5.0 %	77.8 %	12.4 %	2.0 %	2.7 %
— muon 10–20 hits	190	72.1 %	0.86	27.9 %	63.7 %	5.8 %	1.1 %	1.6 %
— muon < 10 hits	194	19.1 %	0.26	80.9 %	13.4 %	4.6 %	1.0 %	0 %
Muonless events	852	30.4 %	0.52	69.6 %	18.2 %	6.7 %	3.6 %	1.9 %
All events with hits	1969	57.3 %	0.79	42.7 %	44.3 %	8.5 %	2.5 %	2.0 %

Detector	Events with a local track	Local tracks / MC particle	Median purity	Local tracks with ≥ 2 MC particles	Median hits / local track
SiTarget (1D strips)	407	1.85	0.80	77 %	15
SiPad (2D pads)	419	1.28	1.00	23 %	8
MTC (SciFi U/ V + Fe)	746	1.76	1.00	10 %	9




Tracking with *ats*: Example of promoted tracks (I)

SiTarget only




Tracking with *ats*: Example of promoted tracks (II)

SiPad only



Tracking with ats: Example of promoted tracks (III)

MTC only




Key4hep **IJCLab**
Irène Joliot-Curie
Laboratoire de Physique
des 2 Infinis

- ▷ Data conversion & Event building
- ▷ Mip & pedestal calibration
- ▷ *Validation variables (shower study)
- ▷ *Event_display (key4hep native) ← *key4ship*
- ▷ *ACTS tracking ← *key4ship*
- ▷ *Event_viewer (Python, real-time tools) – O. Arquero¹
- ▷ *Jupyter notebooks (python, dynamic analysis) – O. Arquero¹

- ▷ Mip Calibration
- ▷ *Realistic segmentation  *key4ship*
- ▷ *Realistic digitization – Y. Shi²  *(soon) key4ship*

- ▷ *Validation variables (shower study)
- ▷ *ACTS tracking ← *key4ship*
- ▷ *Event_display (key4hep native) ← *key4ship*
- ▷ *Event_viewer (Python)




Reusable components across SHiP & TB software → prototype changes propagate efficiently through the full pipeline

Synergies between key4ship & siwecal software (I)

► Realistic segmentation of the SiPad detector has been developed and added in the TB software, and then transferred to key4ship

▷ Real dimensions are $5.52 \times 5.52 \text{ mm}^2$ with 0.01 mm gaps in between pixels

▷ Margin in the outter pixels $3.375 - (5.52 + 0.01)/2 = 0.61 \text{ mm} \rightarrow 1.22 \text{ mm}$ gap



Description from the manufacturer



e- shower from TB data (note the cross-gap in the middle)



- ▶ Key4ship's optimization is built on BoTorch (Python)

- ▶ **Sobol sampling:**

- Sampling a multidimensional space in a well-distributed way
- Much more efficient than a regular grid in high dimensions
- Less clusters and empty regions compared to standard random sampling
- Well suited for the initial exploration of the parameter space

- ▶ **qLogNEHVI** (Noisy Expected Hypervolume Improvement, fine-tuning):

- q — batch: Picks several diverse candidates at once, greedily, so the batch doesn't clump
- Log: Log-space reformulation that fixes vanishing gradients
- N — Noisy: Handles noisy observations by integrating over previously evaluated points too
- E — Expected: Balances good predicted values against uncertainty
- HVI — Hypervolume Improvement: How much a new point grows the volume of objective space dominated by the Pareto front

Exploded view: gray 3D cluster with RGB-coded 2D projections (N = 343 in a 1 m³ cube)

