Autonomous Tracker Design

A systematic search methodology for high-dimensional tracker design

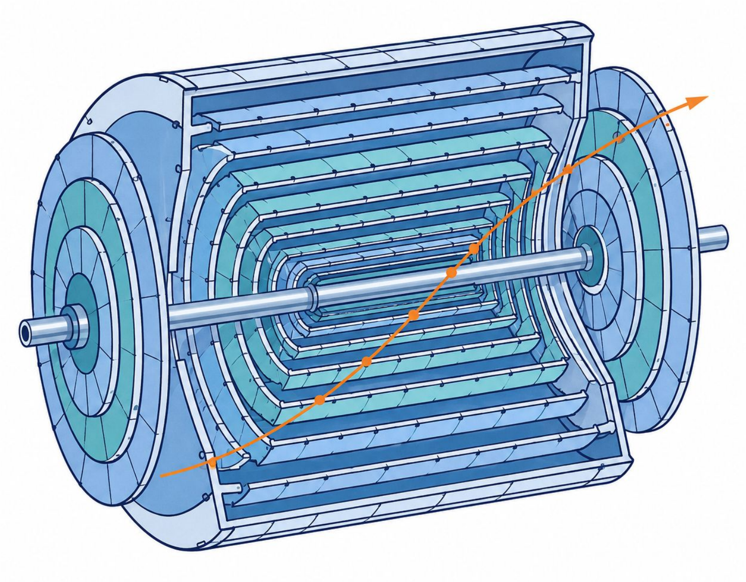

Yiren Wu

National Taiwan University

Team: Yiren Wu, Hanyi Ying,
Olivia Holmberg, Loukas Gouskos

Why a better tracker is hard to find

Every candidate costs a simulation, the count explodes with dimension, and parameters cannot be tuned one at a time

01

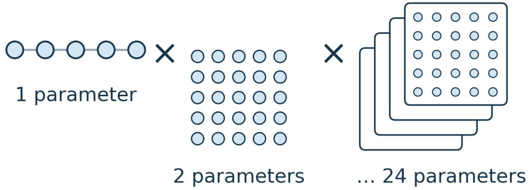

Too many combinations

24 continuous parameters; the number of combinations grows with every one added

02

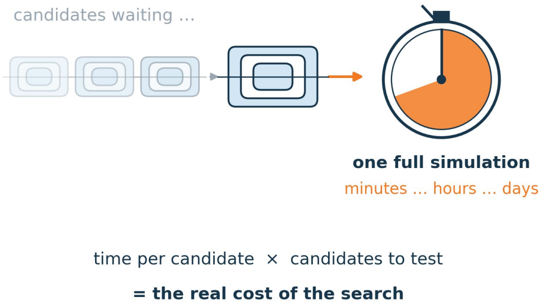

Simulation takes time

Every candidate is one full simulation, so only a small fraction of the space can ever be tested — less still with a realistic simulator

03

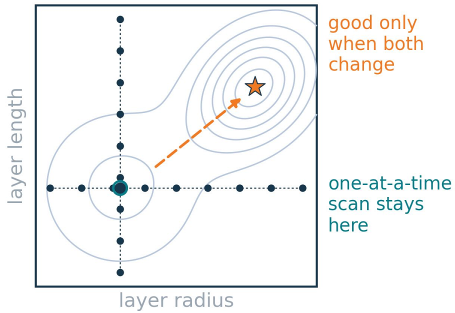

Parameters are coupled

Moving one layer changes what the others contribute; a one-at-a-time scan misses geometries that only work as a combination

Challenge: find useful parameter combinations without scanning the full design space

This study is a starting point for more complex subdetectors, such as calorimeters.

A systematic search for better trackers

Past results tell the framework which geometry is worth simulating next

01

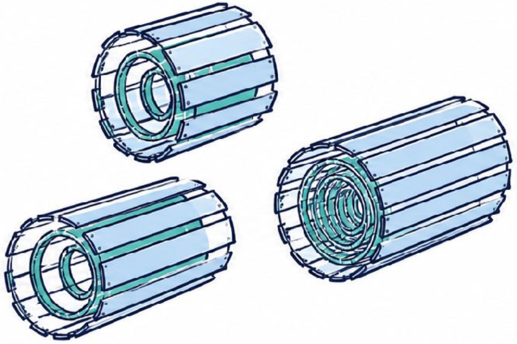

Design space

24 coupled geometry
variables to explore

02

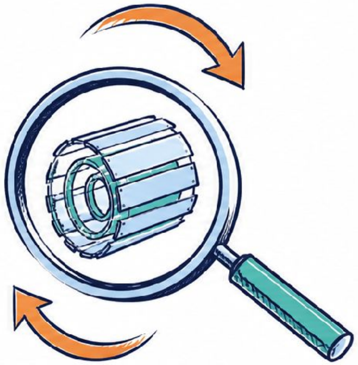

Learn, then propose

Each simulated geometry
updates a model;
the model ranks what to try next

03

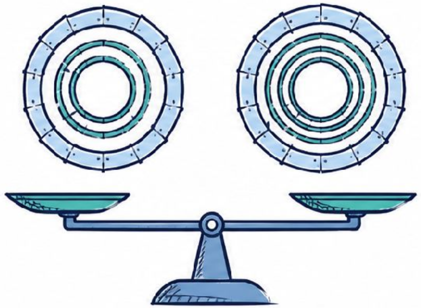

Evaluation

Compare with IDEA
and earlier search methods

Every simulation teaches the framework where the next promising geometry is

IDEA tracker optimization below 60 GeV

FastTrackCovariance evaluates IDEA and every candidate geometry with the same simulation

https://github.com/cvernier/FastTrackCovariance.git

IDEA reference geometry

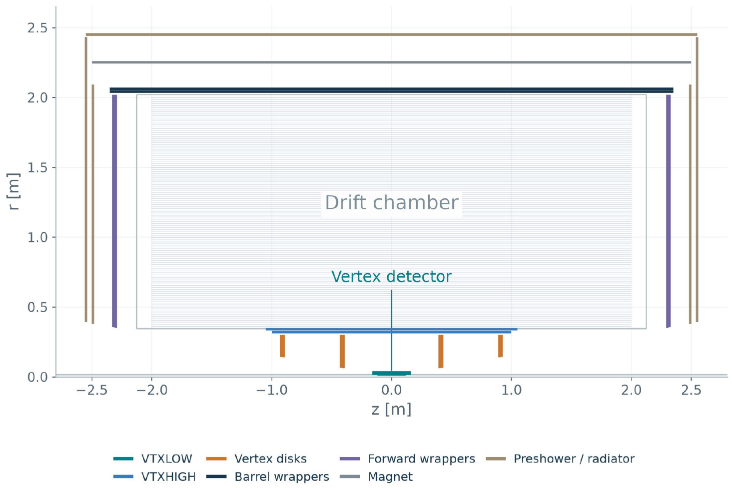

Full detector model, including both z sides

Study range: total momentum p < 60 GeV

Resolution below 60 GeV

θ = 45° B = 2 T

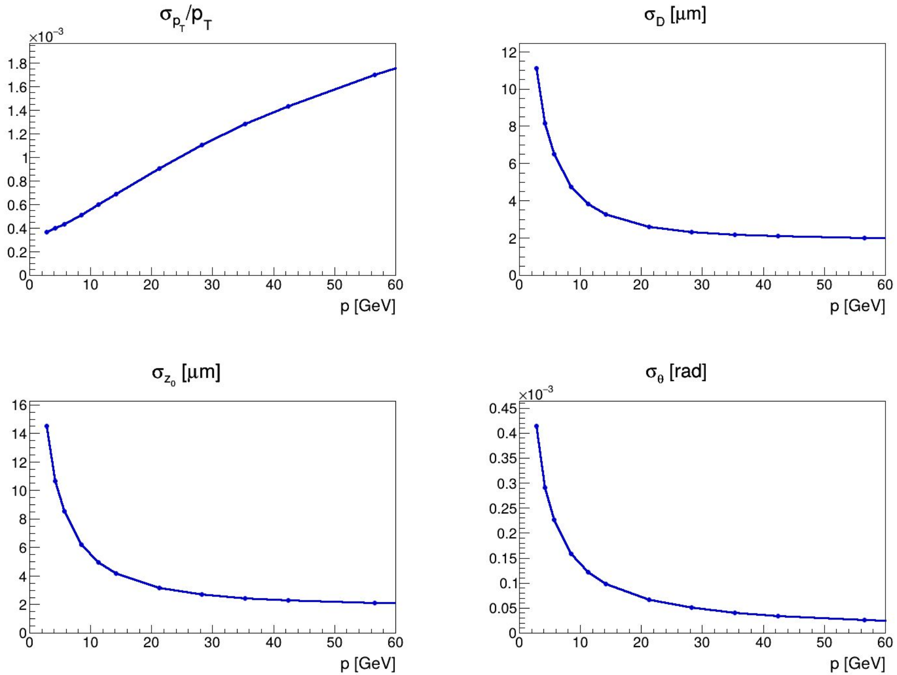

Global GP BO and our method

We compare complete methods under the same physics objective and evaluation budget

Component	Global GP BO control	Our method
Surrogate target	Raw physics score	Rank-Gaussian score
Model structure	One global exact GP	Global GP and local exact GP
Local adaptation	No explicit trust region	Inherited ARD scales and staged fitting
Batch selection	GP-Hedge: LCB / EI / PI	qLogEI inside a TuRBO region

The global model selects a region. TuRBO adapts the local search.

Several components change together. The matched-budget result evaluates the complete method.

Control: our tracker adaptation of dRICH-inspired BO. Cisbani et al., JINST 15 (2020) P05009.

24 coupled variables define the design space

We vary layer dimensions and positions within a fixed tracker topology

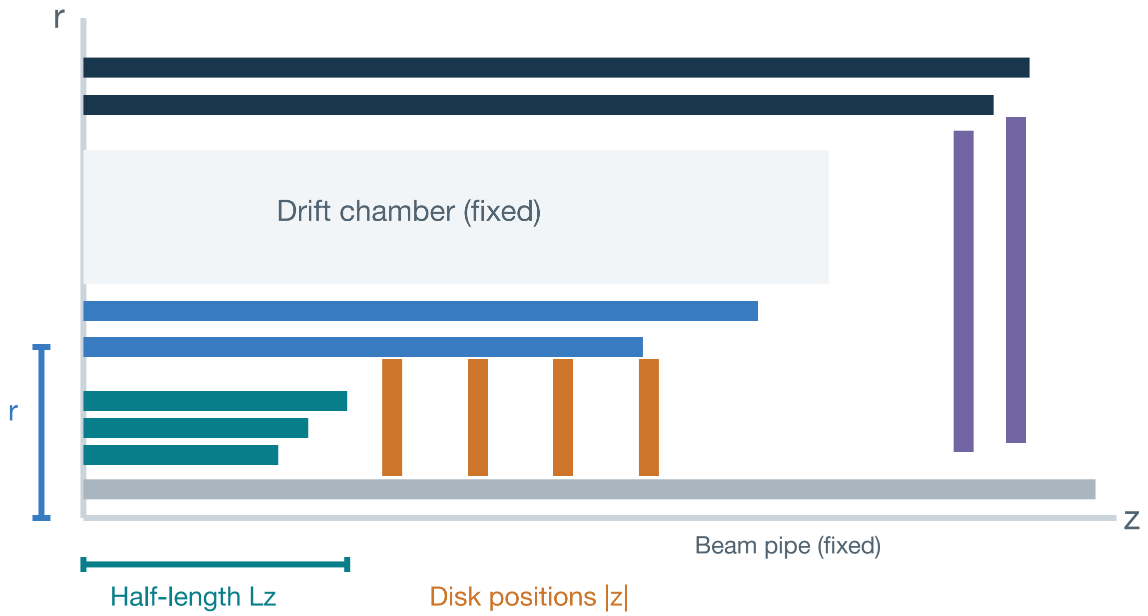

r–z schematic, not to scale; +z half shown

VTXLOW barrels — **6D**
3 × (radius, half-length)

VTXHIGH barrels — **4D**
2 × (radius, half-length)

Barrel wrappers — **4D**
2 × (radius, half-length)

Ordinary disks (LLHH) — **4D**
4 pairs × longitudinal position

Forward wrappers — **6D**
2 pairs × (|z|, ,)

Finding promising combinations requires an efficient search

Bayesian optimization guides the geometry search

Learn from evaluated designs to choose the next simulations

01 GP surrogate

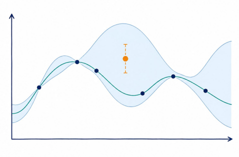

Learn geometry quality and uncertainty

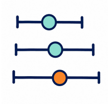

Prediction + uncertainty

02 BO selection

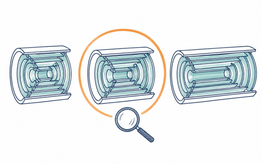

Choose the next batch of designs

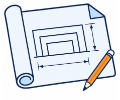

Candidate geometries

03 FastTrackCovariance

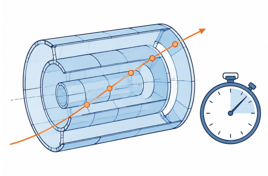

Evaluate tracking resolution

Resolution data

04 Score against IDEA

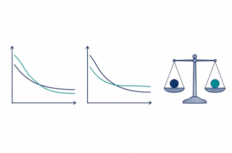

Aggregate resolution improvements

New scores

Repeat with new evidence

Each simulation informs the next geometry choice

Conceptual illustrations

We score each geometry relative to IDEA

A common physics objective for every search method

Four tracking-resolution targets

d_0 **$p_{\square}$** **z_0** **$\cot\theta$**

Evaluate each target over total momentum $p < 60$ GeV

$$s(x) = \Sigma_{\square}\, w_{\square}\, [1 - \sigma_{\square}(x) / \sigma_{\square}(\mathrm{IDEA})]$$

$\Sigma_{\square}\, w_{\square} = 1$

IDEA = 0

Positive score: an overall resolution improvement

The weighted score allows trade-offs between targets. It does not imply improvement at every momentum or angle.

Rank transformation improves surrogate ordering

Learn a scalar surrogate on transformed scores; retain the raw score for physics comparisons

$$r_i = \mathrm{rank}^D(s_i) \quad \rightarrow \quad u_i = (r_i - ½) / N \quad \rightarrow \quad z_i = \Phi^{-1}(u_i)$$

D = GP training subset; N = |D|. Ascending ranks; ties averaged. Φ = standard normal CDF.

Actual local training data

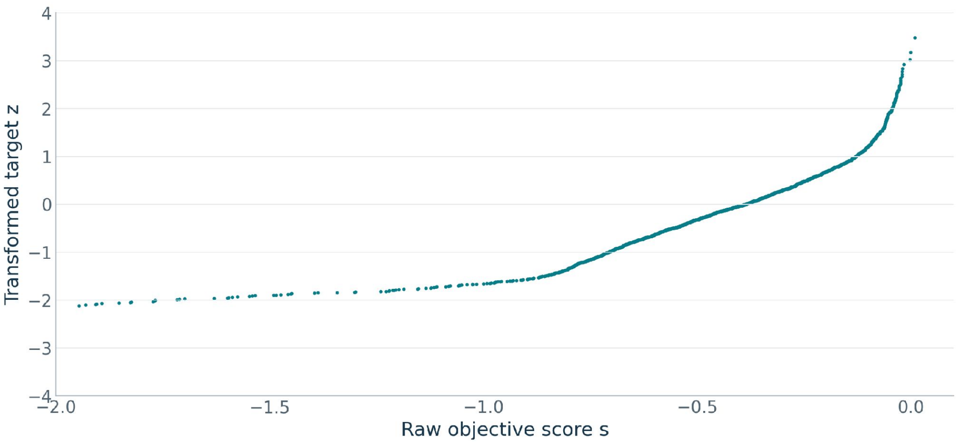

Separate held-out surrogate ablation

15k train / 5k validation · SVGP, M = 512 · frozen kernel

Validation metric	Raw	Rank
Spearman ρ ↑	0.791	0.967
Top-10% recall ↑	37.2%	84.6%

Better ordering

Earlier full-range objective; separate from the p < 60 GeV BO comparison.

GP + qLogEI use z. Incumbent tracking + trust-region feedback use raw s.

Global guidance initializes the local GP

Local fitting becomes more flexible as evaluated designs accumulate

Global Rank-GP

Learn from all evaluated designs at initialization or restart.

Guided initialization

Use nearby evaluations and global ARD scale ratios.

Local GP and TuRBO

Propose each batch within the current trust region.

Few local observations

Fit a shared scale while retaining the relative parameter scales learned by the global model.

More local observations

Fit individual ARD length scales and gradually relax their priors as local evidence grows.

Why stage the fit? With little local data, estimating all 24 scales is difficult.

This is the design rationale. The current comparison does not isolate the benefit of staged fitting.

TuRBO adapts the trust region

The raw physics score determines whether the search expands, contracts or restarts

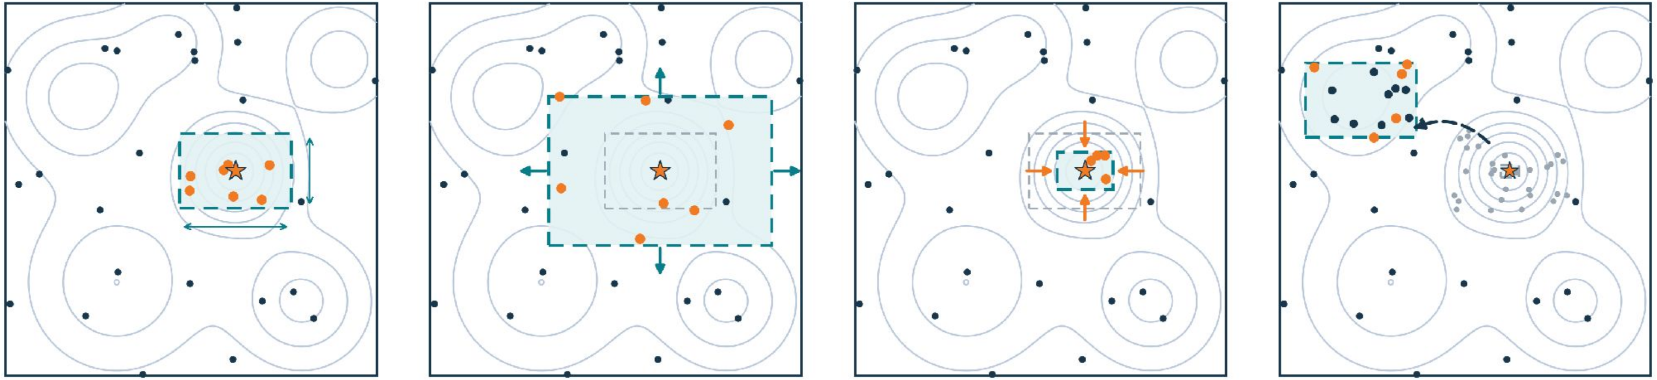

Propose locally

Choose the next batch
inside the region.

Expand on improvement

Repeated gains allow
a larger search region.

Shrink on stalls

No progress narrows
the search region.

Restart

Global guidance selects
another starting region.

The trust region responds to observed raw-score improvement.

Trust-region adaptation in the optimization run

Observed region length in the same run used for the performance comparison

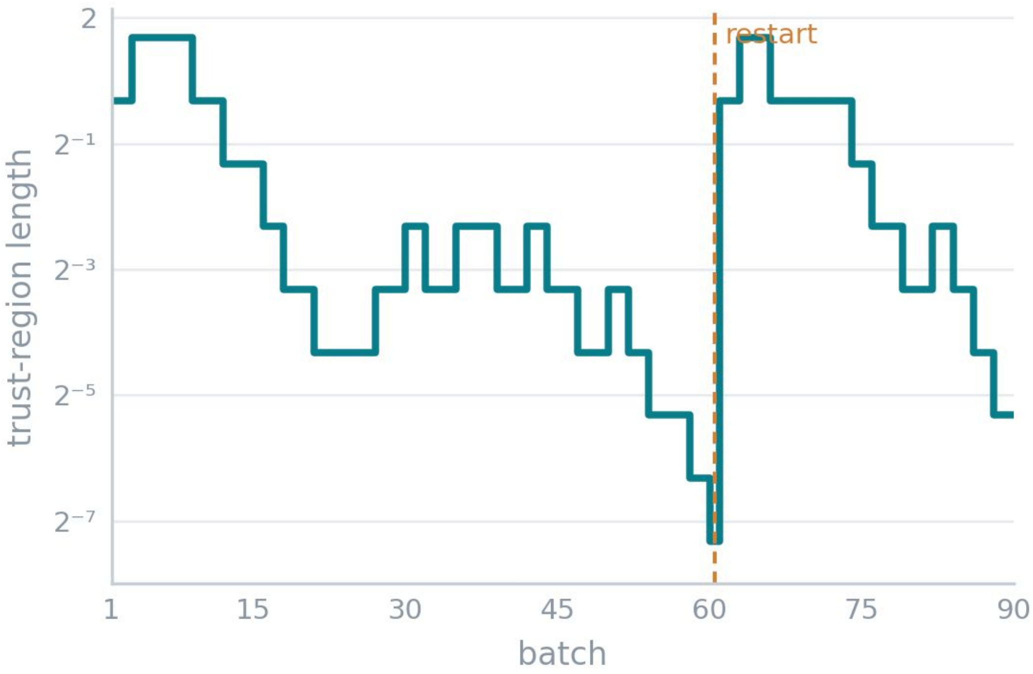

Expansion

Repeated improving batches increase the region length.

Contraction

Stalls reduce the region length and concentrate the search.

Restart at batch 61

At the minimum length, global guidance seeds a new region.

This trace shows how the search operated. The next comparison tests the complete method’s performance.

Higher best score under a matched evaluation budget

Same 200 initial designs · 20 evaluations per batch · 2,000 total evaluations

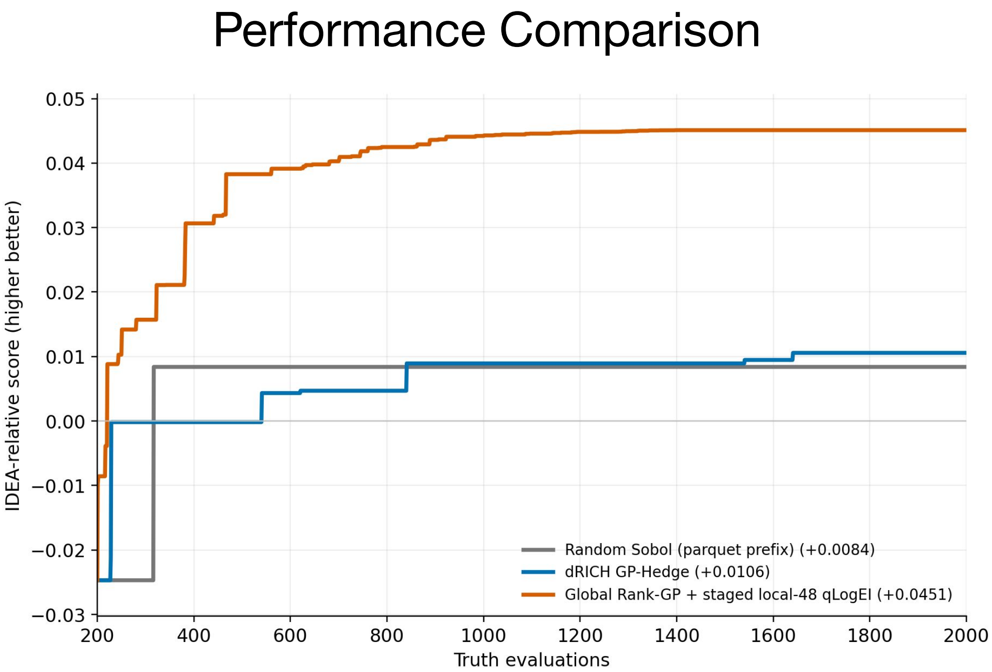

Best designs vs IDEA

Weighted resolution margins

Target	Control	Ours
d_0	–0.09%	0.79%
σ(p□) / p□	0.50%	6.67%
z_0	1.21%	3.09%
cot θ	2.61%	7.49%
Overall s	1.06%	4.51%

IDEA = 0

The complete method reaches a higher weighted score in this matched run.

One run per method. Aggregate gains do not imply improvement at every momentum / angle point.

Best geometry vs. IDEA

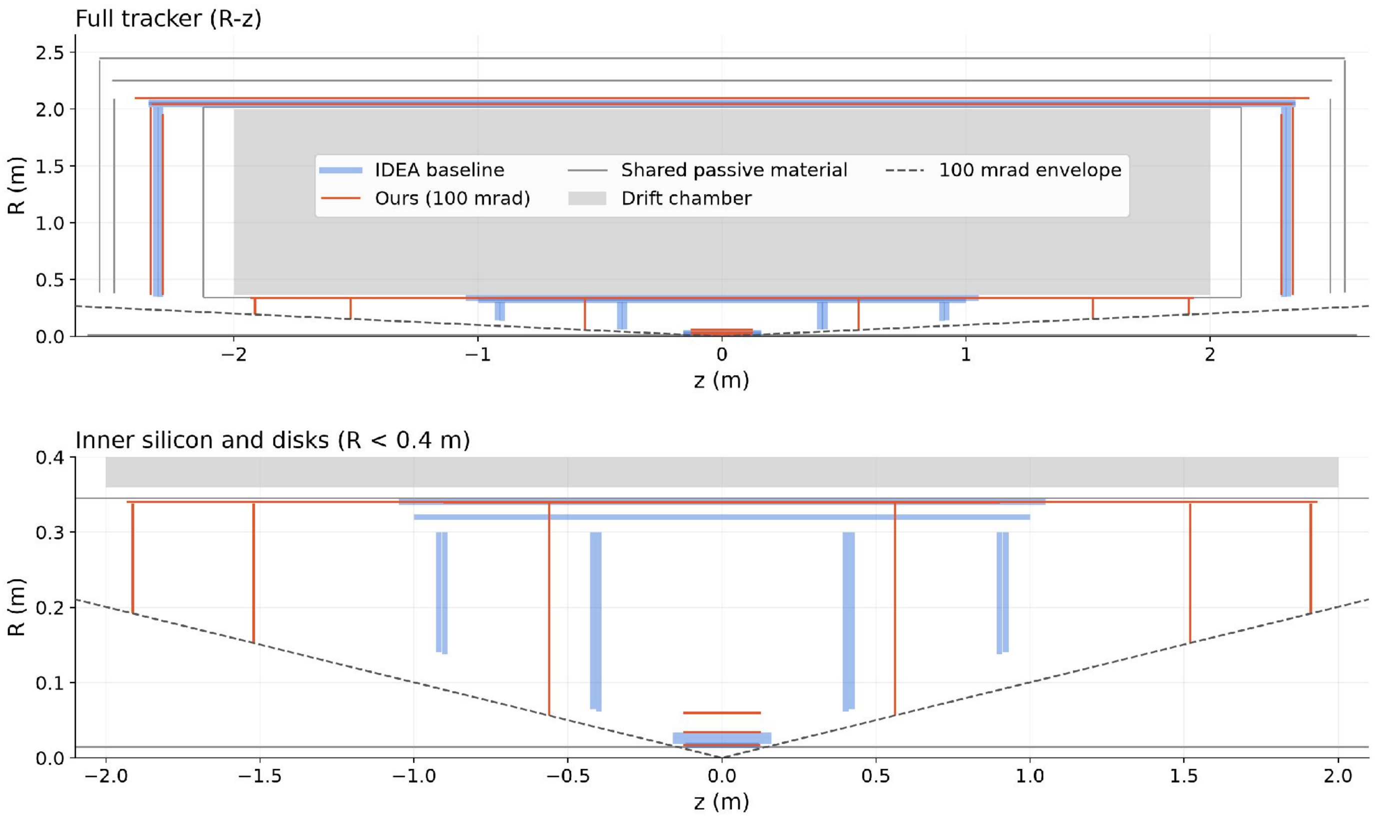

The correction leaves evaluated resolutions unchanged. Detailed beam-pipe, support and cooling material remain unmodeled.

Optimized geometries: ours and global BO

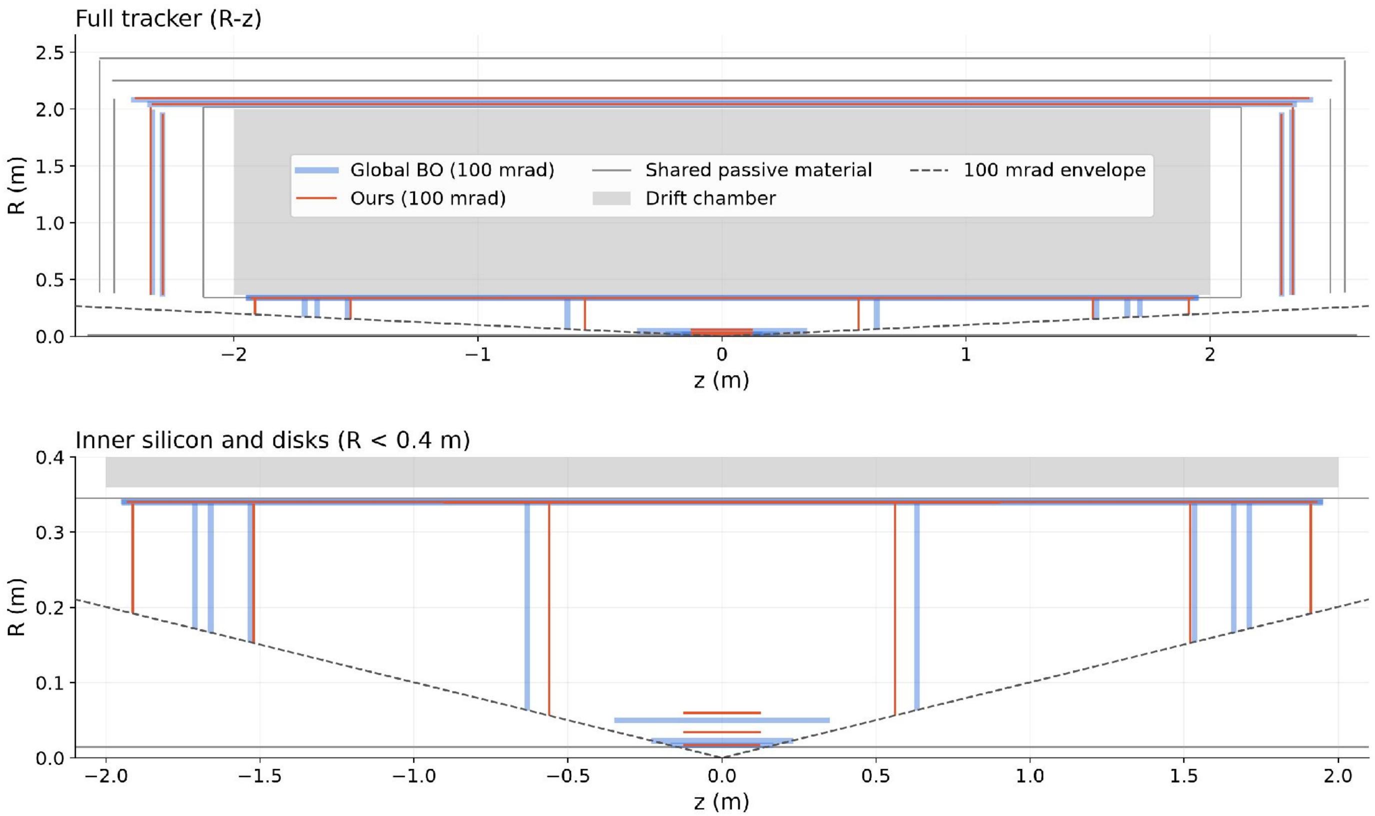

Both geometries satisfy a simplified 100 mrad envelope. Detailed engineering remains outside this model.

IDEA, global BO and ours at θ = 10°

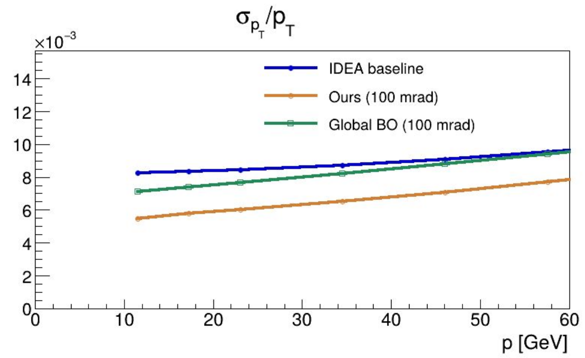

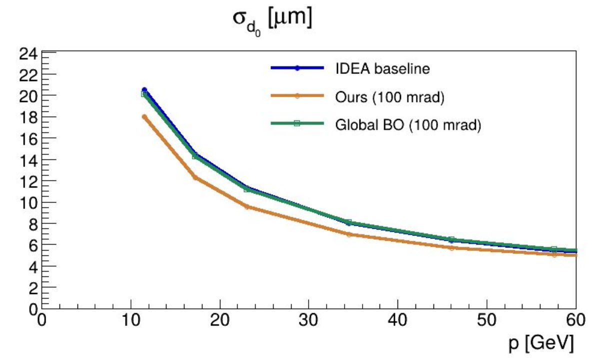

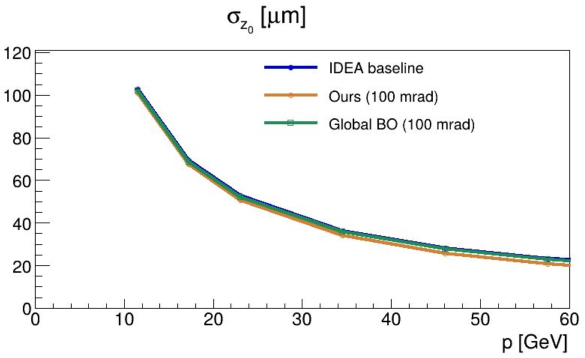

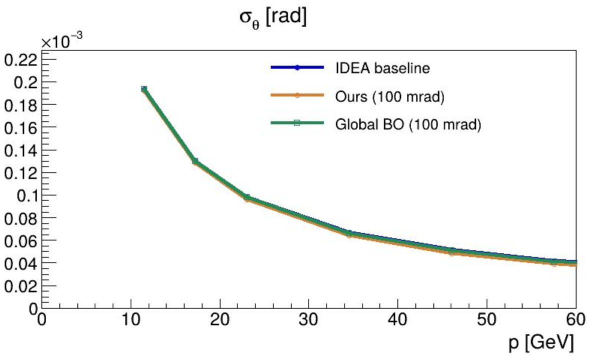

Resolutions unchanged after correction. Samples start at p = 11.52 GeV.

IDEA, zglobal BO and ours at θ = 20°

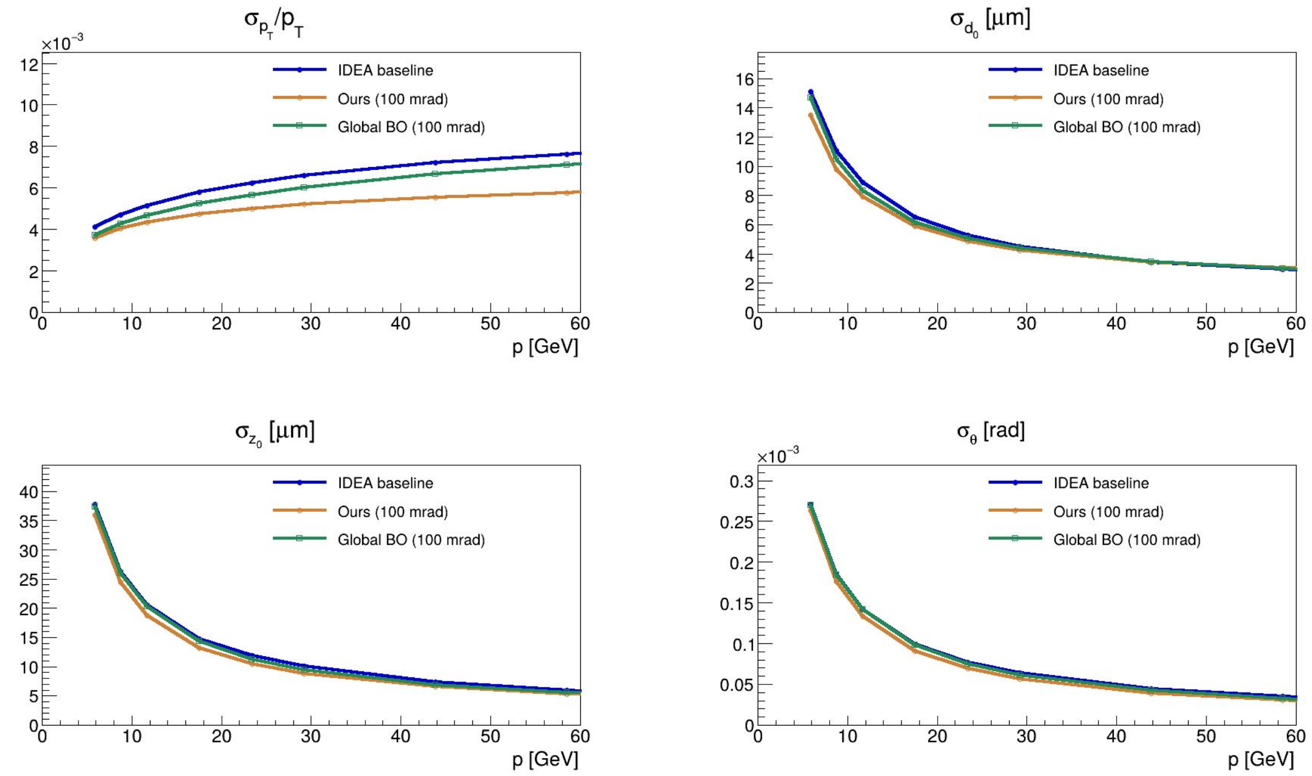

Resolutions unchanged after correction. Samples zstart at p = 5.85 GeV.

IDEA, global BO and ours at θ = 60°

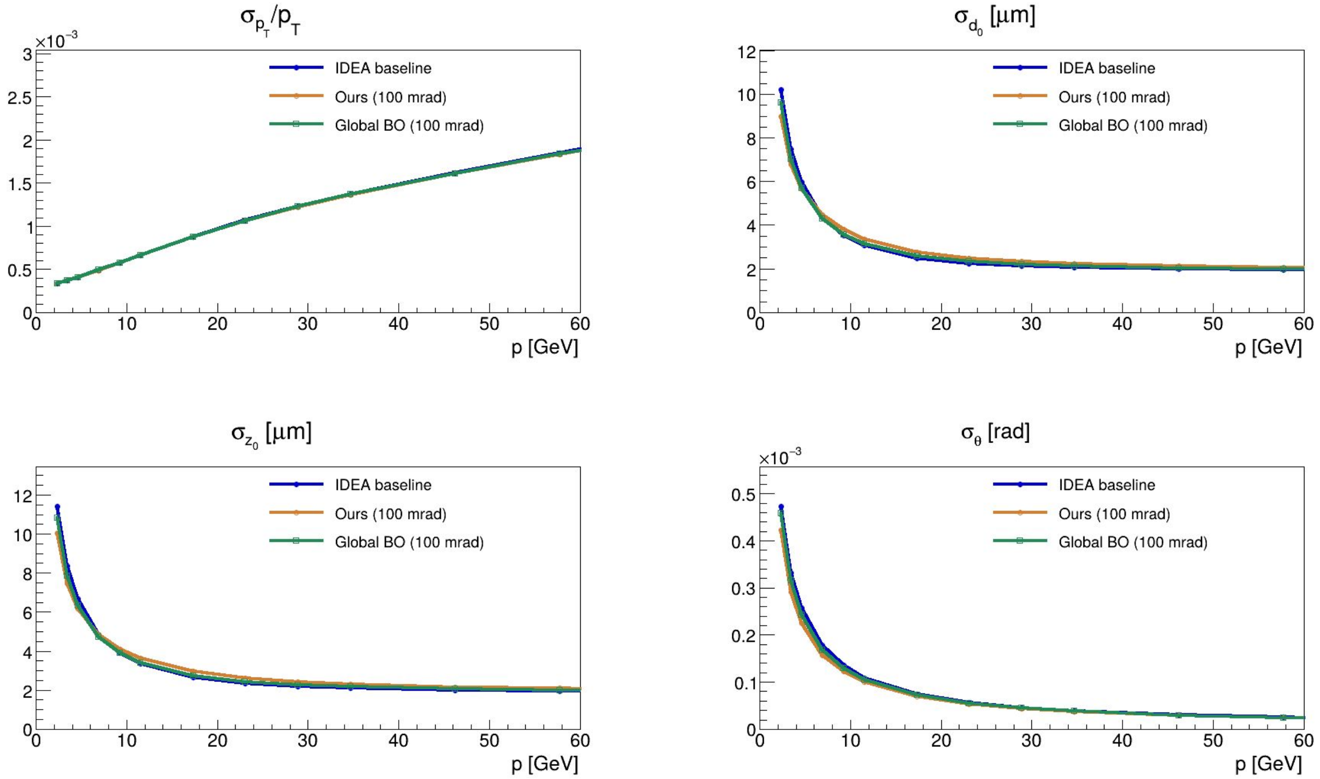

Resolutions unchanged after correction. Samples start at p = 2.31 GeV.

Detector optimization for different physics goals

Our IDEA study demonstrates the search procedure for one objective and a defined design space.

Different objectives

Study how score functions aligned with different physics targets change the geometries we discover. Vary resolution weights and priorities across momentum and angle.

A broader design space

Allow more geometry parameters to vary and test whether greater design freedom leads to further improvements.

Future studies: how do the preferred geometries change with our goals and design freedom?