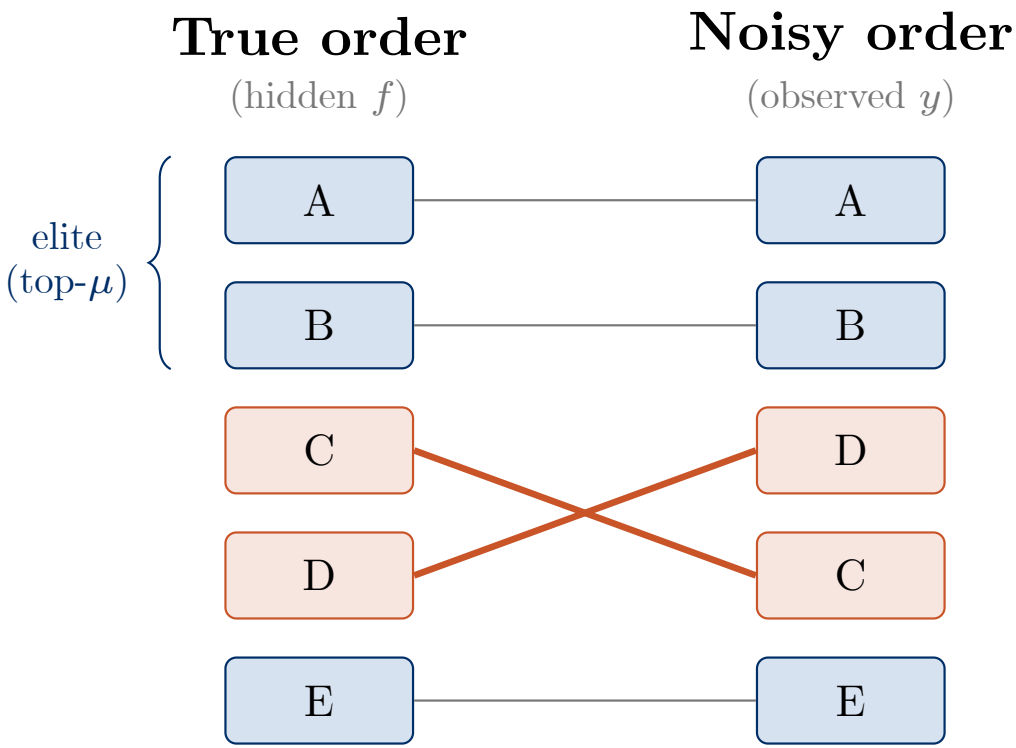
Depth over Fidelity
in Fixed-Budget Noisy Evolution Strategies

Sichen Wang Zhipeng Lu

Shenzhen MSU-BIT University

ICML 2026

The Problem: Misranking



- CMA-ES uses only the **ranking** of candidates.
- Noise flips ranks — a **misranking**.
- The wrong candidate enters the elite set, **corrupting the update**.

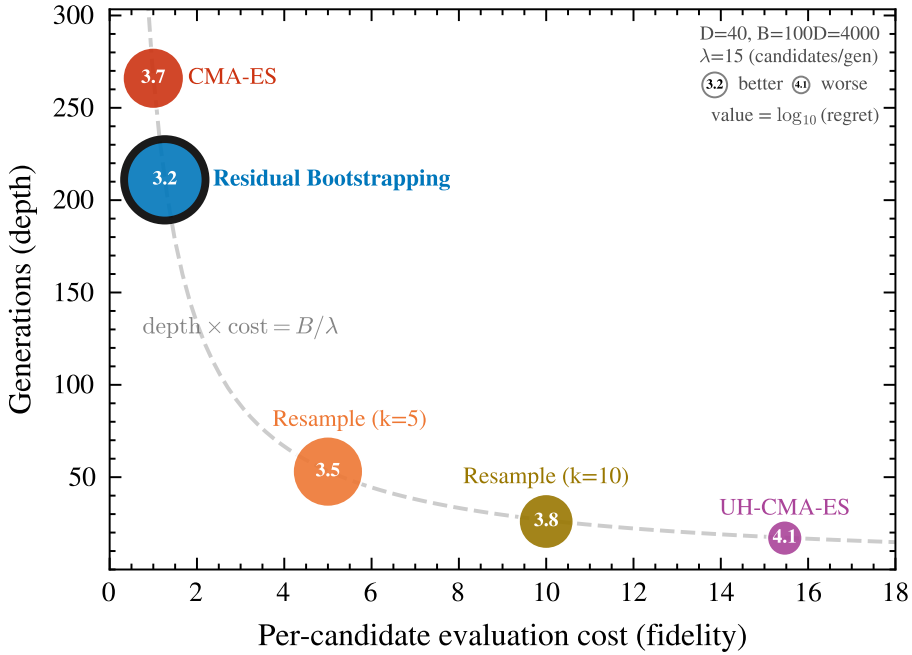
The Trade-off: Fidelity Costs Depth

- Each generation samples λ candidates; resampling each k times denoises the ranking — higher **fidelity**.
- But under a fixed budget, depth shrinks:

$$T \approx \frac{B}{k\lambda}.$$

(λ = population size)

- So fidelity is bought with **depth** — the number of update steps.



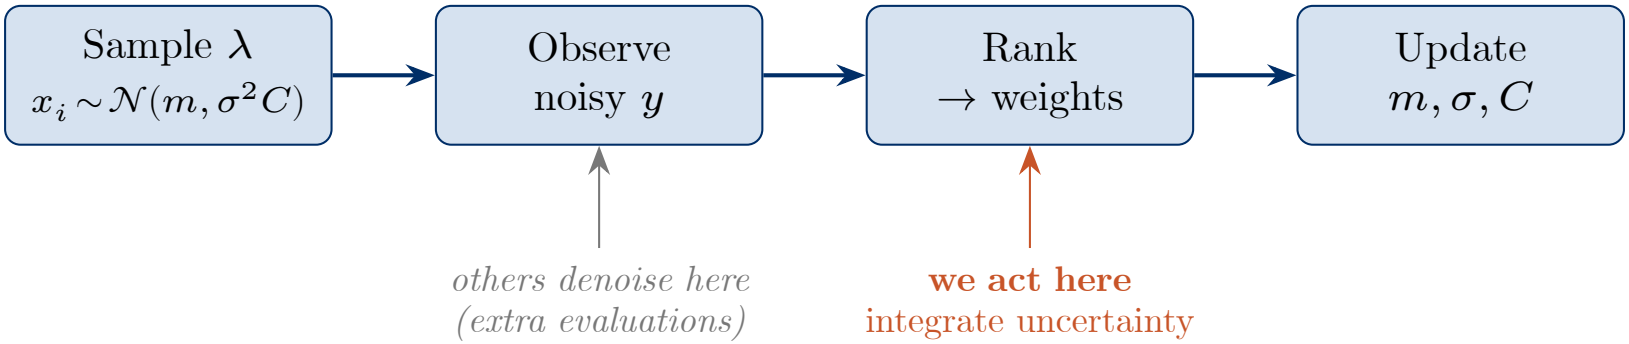
Methods on the budget frontier; denoising → fewer generations.

Depth over Fidelity

When the ranking is noisy,
don't pay to clean it up.

Absorb the uncertainty into the update
— and keep iterating.

Where We Intervene

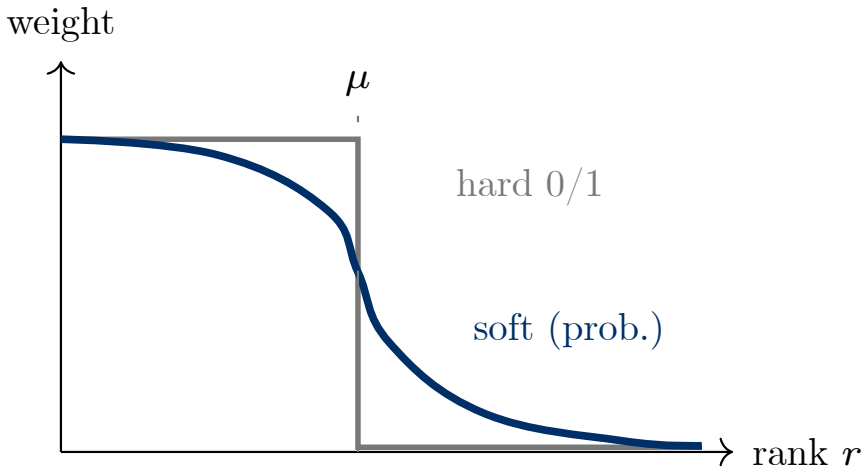


$$m^+ \;=\; m + \eta \sum_i w(r_i)\,(x_i - m)$$

CMA-ES mean update — PEM will replace the weights $w(r_i)$

Keep a **one-evaluation baseline** — with only small capped overhead, we change how ranking becomes weights $w(r_i)$.

PEM: Probabilistic Elite Membership



- Replace the hard elite label by a **probability**:

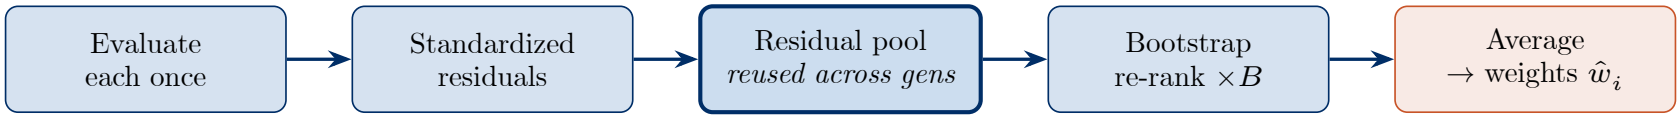
$$w_i^\star = \Pr(r_i \le \mu \mid x_{1:\lambda}) \,.$$

- A borderline candidate gets a *fractional vote*.
- A **Rao–Blackwellization** of the noisy step:

$$\Delta m_{\mathrm{PEM}} = \mathbb{E}[\Delta m^{\mathrm{noisy}} \mid x_{1:\lambda}]$$

same expected step, less variance.

RB-PEM: Residual Bootstrapping

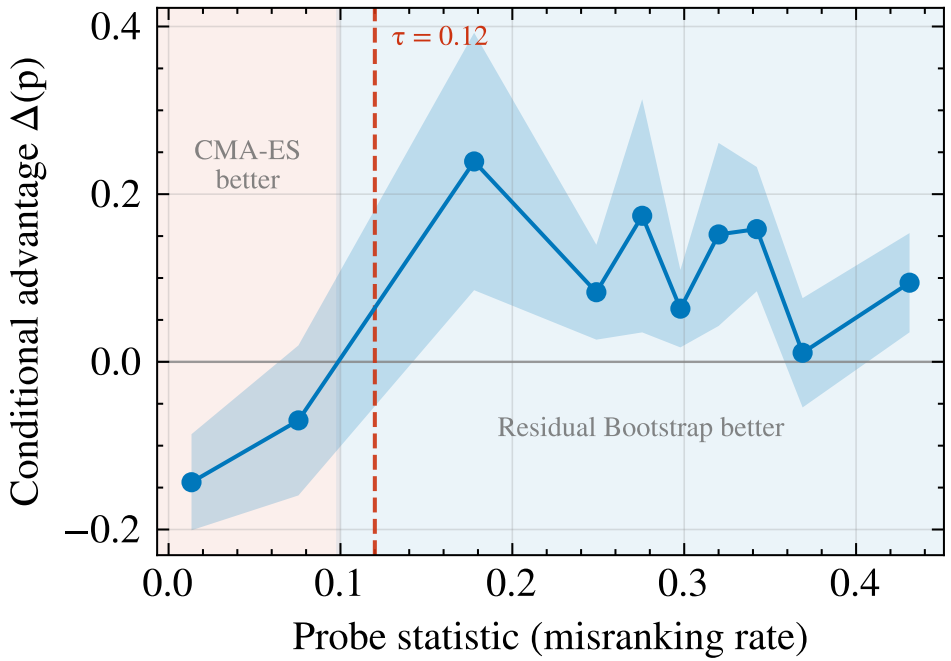


$$\hat{w}_i = \frac{1}{B}\sum_{b=1}^{B} w\big(\mathrm{rank}_i(\tilde{y}^{(b)})\big), \qquad \tilde{y}^{(b)} = \hat{f}(x) + \hat{s}(x)\, z^{(b)}, \quad z^{(b)} \sim \widehat{D}$$

synthetic values from the fitted noise model $\widehat{D}$ — re-ranked, then averaged

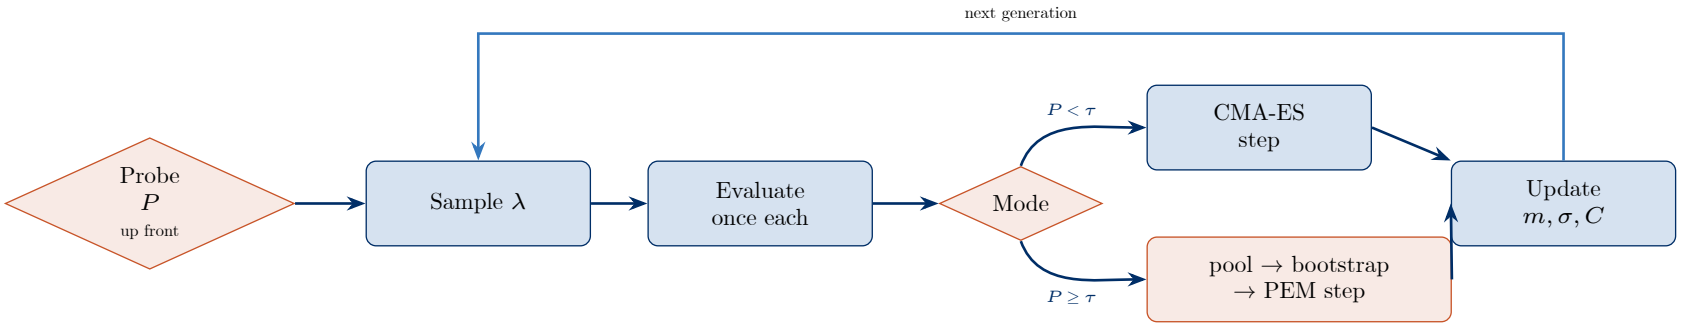
Bootstrap rankings are **software-only**; targeted re-evaluation is capped at $K_{\max} \leq 1$, so depth stays near B/λ.

Probe-and-Switch



- A cheap **probe** P measures rank instability.
- $P \geq \tau \Rightarrow$ run RB-PEM.
- $P < \tau \Rightarrow$ fall back to **CMA-ES**.
- **Robust**: gains where they exist; reverts when ranks are stable.

The Full Algorithm



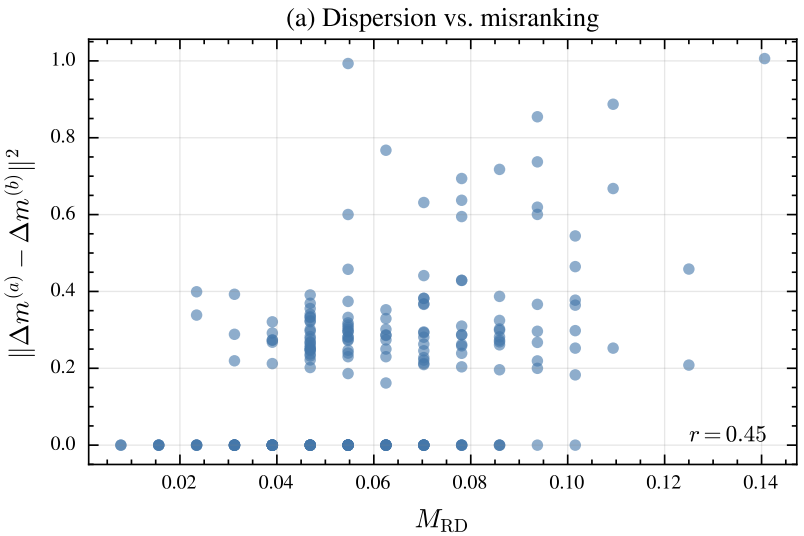
Probe once, then use a one-evaluation baseline plus capped overhead when active — **depth is near-maximal**.

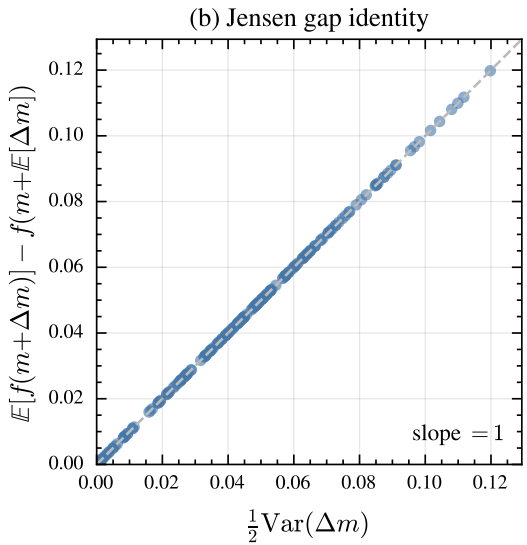
Why It Works: The Curvature Tax

$$\mathbb{E}[f(X) \mid x_{1:\lambda}] \;\geq\; f(\bar{X}) \;+\; \frac{\alpha}{2}\,\mathbb{E}\Big[\,\|X-\bar{X}\|^2 \,\Big|\, x_{1:\lambda}\Big]$$

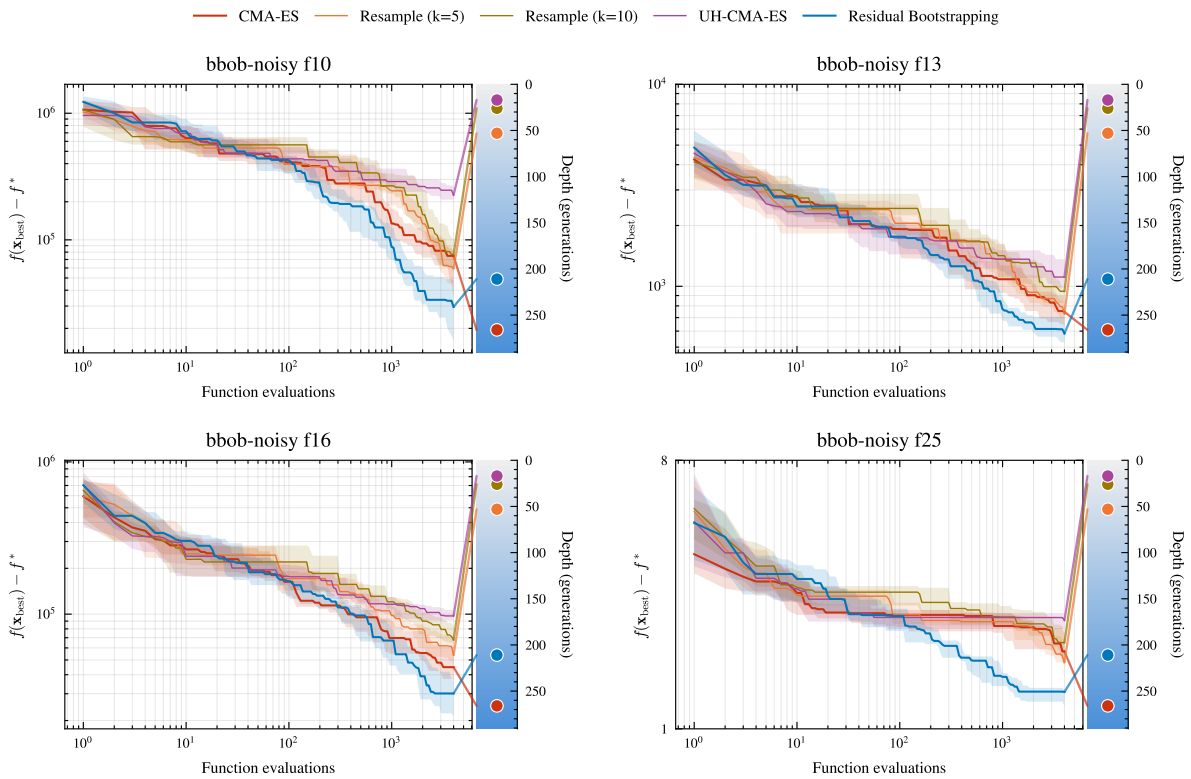
Theorem 1 (proof in paper) α: local curvature $\|X-\bar{X}\|^2$: update dispersion

Update **dispersion** $\times$ curvature = an unavoidable per-generation tax. PEM minimizes the variance, so it minimizes the tax.





Evidence



Probe-and-Switch win rate (COCO + 6 external tasks):

vs Res.(10)	vs Res.(5)	vs UH-CMA-ES	vs CMA-ES
85.6%	**79.3%**	**77.4%**	54.3%

Ranks strongest on high-misranking functions, ahead of LRA-CMA-ES; the CMA-ES edge grows with dimension.

Takeaway

Depth over Fidelity

- Under noisy rankings + tight budgets, **integrate** uncertainty into selection rather than spending evaluations to remove it.
- Largest gains where misranking is high; *robust* elsewhere.
- Portable to other population-based optimizers.

Thank you!

github.com/sichen-wang/Depth-over-Fidelity_ICML2026

Code