How Many Initial Points Does Bayesian Optimization Need?

Anonymous Author(s)
Affiliation
Address
email

Abstract

Bayesian Optimization (BO) pipelines generally begin with an *initialization phase*: a batch of n_0 uninformed evaluations. The choice of n_0 is generally guided by rules of thumb, and we empirically observe that the total cost (random initial points plus BO iterations needed to find the global optimum) is U-shaped in n_0, i.e., a practitioner will waste resources by selecting either too low or too high a value of n_0. We find this tradeoff persists across MLE, Bayesian MCMC, and *exact* GP hyperparameters, as well as across acquisition functions. Toward the latter, Thompson Sampling appears the exception, with both total cost and simple regret essentially n_0-agnostic, though at a uniformly higher level. We attribute this U-shape to the known boundary issue of variance-driven BO: BO burns early budget on corners of the hypercube before turning inward. We demonstrate this effect using a 3D BO trajectory where the exact hyperparameters are known. We conclude with some practical recommendations: using Thompson Sampling for n_0-insensitive deployments, and taking a generously large n_0 to improve downstream BO performance.

1 Introduction

Sequential decision-making with expensive evaluations—whether in Bayesian Optimization (BO) for molecular design [Xie et al., 2025], contextual bandits for recommendation [Li et al., 2010], or engineering systems design [Cheon et al., 2024, Paulson and Tsay, 2025]—routinely combines an offline phase of uninformed exploration with online, model-guided adaptation. In BO [Jones et al., 1998, Frazier, 2018, Garnett, 2023], this transition is ingrained in the standard protocol: evaluate n_0 initial points (by random, Latin-hypercube, or Sobol sampling), fit a Gaussian process (GP) surrogate, and then transition to using an acquisition function to guide samples. The first phase is effectively an offline dataset (and ignored in most BO literature); the second phase, where BO takes over, is online adaptation. Despite the ubiquity of this two-phase protocol, the question of how to size the offline phase has received relatively little attention. Practitioners often rely on simple rules of thumb; for instance, Meta's Ax platform [Olson et al., 2025] defaults to 5 initial trials (the search-space center plus 4 quasi-random Sobol points), independent of dimensionality. Another common rule for the GP modeling is $n_0 = 10d$ [Loeppky et al., 2009], no analogous analysis exists for *total experiment* cost.

The initialization tradeoff. Let $\tau(n_0)$ be the number of BO iterations needed to find the global optimum after n_0 random initial evaluations and $C(n_0) = n_0 + \mathbb{E}[\tau(n_0)]$ the total expected cost. We empirically observe that $C(n_0)$ is sharply U-shaped on discrete grids, and that this U-shape persists across (i) hyperparameter regimes (MLE / Bayesian MCMC / oracle) and (ii) multiple variance-driven acquisition functions. Toward the latter, we observe that the total cost and convergence of BO with EI, UCB, and PI is sensitive to the choice of n_0, while Thompson Sampling [Thompson, 1933], or TS, appears the lone exception (Section 2). We connect the U-shape to the known *boundary issue* of variance-driven BO [Swersky, 2017, Oh et al., 2018, Siivola et al., 2017]: GP posterior variance peaks at points farthest from observed data, and, for a box-shaped domain, those farthest points

are the corners (and, in higher dimensions, the faces and edges). Prior work mitigates this through architectural remedies (cylindrical kernels, virtual derivatives, trust regions), we show that the choice of n_0 alone meaningfully shortcuts the boundary itinerary. We demonstrate the trap directly on a 3D Oracle BO trajectory (Section 3) show that random initialization can mitigate these inefficient samples in practice.

Contributions. This work documents an initialization tradeoff in the BO total cost $C(n_0)$ on discrete grids across various hyperparameter settings and acquisition functions, with Thompson Sampling as a notable n_0-agnostic exception. We also present experiments indicating that a similar trade-off exists in the continuous variable setting. We connect this tradeoff to the known *boundary issue* of variance-driven BO and formalize its mechanism as a corner-by-corner itinerary, demonstrated directly on a 3D Oracle BO trajectory (Section 3). Whereas prior work addresses the boundary issue at the level of the surrogate or optimization scheme, we identify n_0 sizing as a simple, complementary lever, and translate this into practical recommendations (Section 4): use Thompson Sampling when one cannot tune n_0, otherwise choose a generously large n_0.

2 The Initialization Tradeoff in Practice

Experimental setup. We consider three hyperparameter settings: *oracle*, where the GP surrogate uses the exact generative kernel, standard *MLE*, and Bayesian *MCMC*. We investigate $C(n_0)$ across three complementary configurations (full details in Appendix A.1): **(E1)** a fixed oracle GP with Matérn-5/2 sampled on the discrete 3D grid $\{-5, \ldots, 5\}^3$ using EI (10^5 episodes per n_0); **(E2)** Ackley 4D on a 9^4=6561 grid with an RBF kernel fitted by either MLE or Bayesian NUTS-MCMC, using EI, UCB, PI, and TS (5,000 episodes); **(E3)** the (E1) Oracle experiment extended to four acquisition functions (EI, UCB, PI, TS), 500 episodes per $(n_0, \text{acquisition})$ pair, $n_0 \in \{0, 10, 20, 30, 40, 50\}$.

The initialization tradeoff is structural. On Ackley 4D (E2), BO with *either* MLE-fitted and Bayesian MCMC-integrated kernels produces sharp U-shapes in $C(n_0)$ (Figures 1–2). Interestingly, while EI, UCB, and PI show acquisition-specific gaps under MLE, the three coincide closely under MCMC. In other words, accounting for hyperparameter uncertainty may mitigate differences between acquisition functions, but the U-shaped initialization tradeoff remains. We further test an oracle GP in (E1), leaving no mismatch between the surrogate model and underlying function, yet the trade-off still appears (Appendix Figure 5). The tradeoff is therefore a structural property of the *decision rule*, and not a result of surrogate model error(s).

Thompson Sampling (TS) is the lone exception. Figures 1–2 reveal a dichotomy: EI, UCB, and PI all exhibit U-shaped $C(n_0)$ under both MLE and MCMC, whereas *Thompson Sampling is the lone exception*, with $C(n_0)$ that is essentially n_0-agnostic at a notably higher absolute level. Experiment (E3) corroborates this regret-side (Figure 3; alternative plots in Appendix A.2), and the same pattern persists on a continuous 3D benchmark (Appendix A.3). Section 3 attributes both to a corner-by-corner boundary itinerary that random initialization shortcuts.

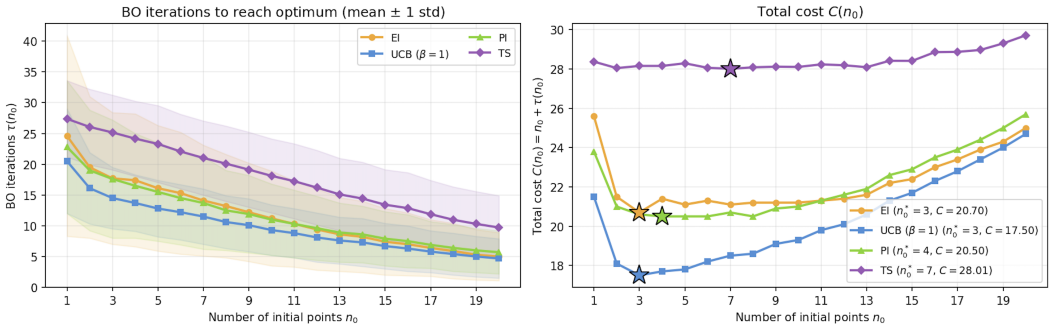

Figure 1: **(E2) MLE-fitted RBF on Ackley 4D.** Total cost $C(n_0)$ over 5,000 episodes (mean $\pm$ std). This is our starting observation: the initialization tradeoff appears strongly in the most widely deployed BO setting.

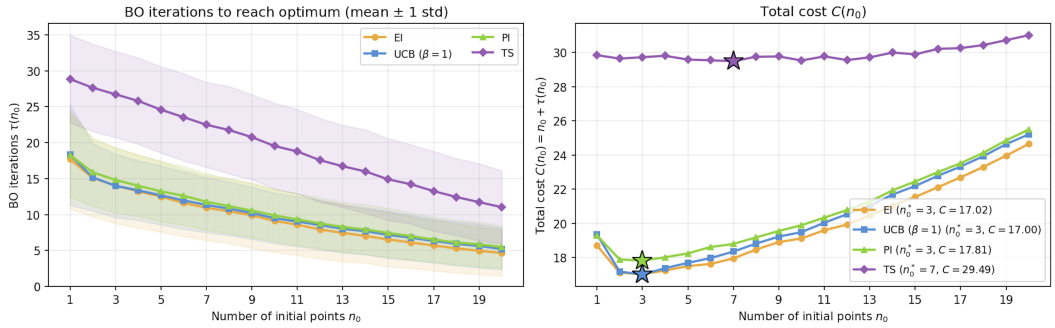

Figure 2: **(E2) Bayesian MCMC RBF on Ackley 4D.** Same setup as Figure 1 but with full kernel hyperparameter posterior. The performance gap among EI, UCB, PI collapses, but the U-shape itself stays. TS remains n_0-flat at a much higher cost.

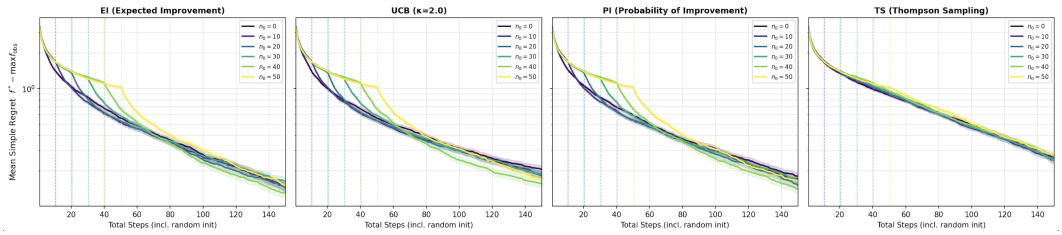

Figure 3: **(E3) Oracle 3D grid, four acquisitions.** Simple regret vs total steps on the (E1) grid $\{-5, \ldots, 5\}^3$, one panel per acquisition and one curve per $n_0 \in \{0, 10, 20, 30, 40, 50\}$ (viridis; dashed = end of random init). TS (right panel) is essentially n_0-agnostic at a higher absolute regret.

3 Mechanism: Boundary Exploration Pathology

Background: the boundary issue. The boundary issue of variance-driven BO is well documented. Swersky [2017] observes that GP posterior variance peaks at points farthest from observed data, which is likely to be on a boundary. This effect diminishes as more points are observed, which Swersky [2017] briefly links to initial-point density.

Refinement. Using variance-driven acquisition functions on a hypercube $[L, U]^d$ can systematically prioritize corner and face regions before the interior, a preference also underlying classical space-filling designs [Johnson et al., 1990]. Consider using UCB with exploration parameter β_t:

$$x_{n_0+1}^{\mathrm{UCB}} = \arg\max_{x \in \mathcal{X}} \left[\mu_{n_0}(x) + \beta_{n_0}^{1/2} \sigma_{n_0}(x)\right].$$

In the low-data regime, the spatial variation of $\beta_{n_0}^{1/2}\sigma_{n_0}$ likely dominates that of μ_{n_0}, and the argmax of UCB coincides with the argmax of σ_{n_0}. EI can exhibit the same variance-dominance: in the limit $\sigma_{n_0}(x) \to \infty$ at fixed $\mu_{n_0}(x) - f_{n_0}^*$, EI reduces to a monotone function of $\sigma_{n_0}(x)$.

Boundary exploration is not inherently wasteful: vertices are directions of maximum posterior variance and could in principle yield important information. What makes vertex-querying pathological is the interaction between (i) the acquisition function's preference for high-variance points and (ii) the kernel's locality. Specifically, under stationary kernels with lengthscale $\ell \ll R$, an observation at a vertex $\boldsymbol{v}$ reduces posterior variance only within an $O(\ell)$-ball around $\boldsymbol{v}$. Therefore, sampling the boundary consumes a number of queries, scaling with the boundary complexity of $\mathcal{X}$, while leaving the bulk of the search domain essentially unexplored.

Random initialization shortcuts this itinerary: scattered initial points reduce the posterior variance everywhere (including near the corners) without the BO loop having to spend queries there. Thompson Sampling escapes the trap by construction: each step samples an independent posterior draw and selects its argmax, which is *typically interior* rather than at the geometric extremes.

Direct evidence (3D Oracle EI, n_0=1). Figure 4 depicts a final-frame snapshot of three independent EI-based BO procedures in setting (E1) with n_0=1 (same fixed f, $x^\star$ in the interior, three different

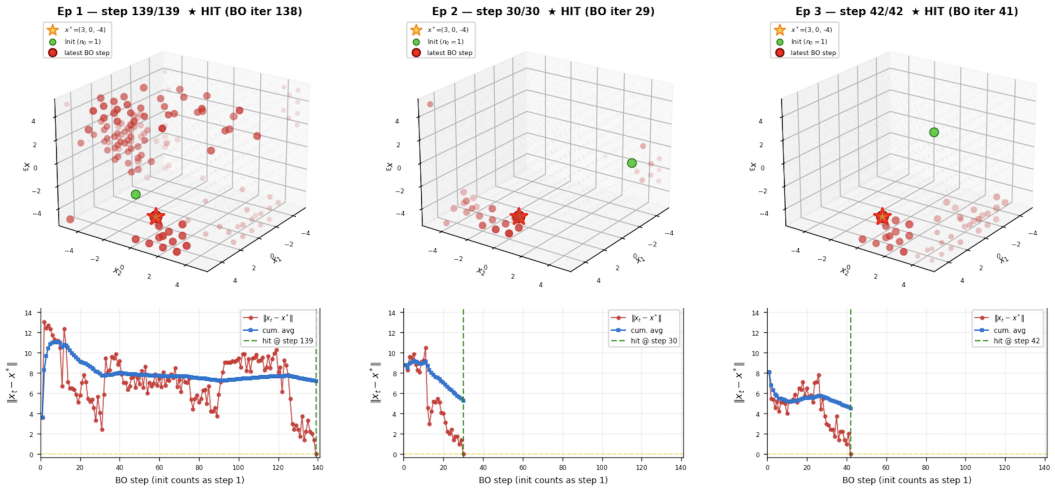

Figure 4: **Direct evidence for the boundary pathology (final-frame snapshot).** Three independent BO episodes on (E1) with n_0=1. *Top row:* cumulative queries (init = green; BO trajectory = fading red, latest = solid red; $x^\star$ = gold star). *Bottom row:* distance $\|x_t - x^\star\|$ vs BO step.

seeds). Figure 4 shows that BO tends to spend its early budget querying the corners farthest from the single initial observation before turning inward, even though $x^\star$ sits well in the interior. With only n_0=1 random observation, the surrogate updated little from the prior, so the variance-driven boundary itinerary is exposed in its worst-case form. An unlucky initial observation pays a large boundary-exploration cost. Moderately larger n_0 shortcuts this itinerary, explaining the left arm of the U-shaped initialization tradeoff; eventually, however, a large n_0 uses too many uninformed evaluations and $C(n_0)$ rises again.

Discussion. This mechanism explains both halves of the dichotomy of Section 2: BO with variance-driven acquisitions over-explores boundaries early, giving rise to a performance trade-off in the selection of n_0. On the other hand, the random sampling of TS breaks the pattern, rendering it relatively agnostic to the choice of n_0. The cost of this independence is that the randomized TS argmax also misses the high-quality argmax in regimes where EI/UCB/PI work well, explaining its higher absolute regret. We do not prove these trends formally; rather, the corner-by-corner itinerary view, together with its link to n_0, organizes the empirical observations of Section 2 into a single causal story.

4 Practical Recommendations and Future Work

When you cannot tune n_0, use Thompson Sampling. Our results show that TS's $C(n_0)$ (Figures 1–2) and simple-regret curves (Figure 3) are both essentially independent of n_0. For deployments that cannot afford to tune n_0 ahead of time, this can be a favorable trade.

Otherwise, prefer a generously large n_0. The Oracle 3D U-shape is markedly asymmetric: $C(1)\approx 94$ and $C(35)\approx 81$, while $C(n_0^\star{=}16)\approx 76.4$ (Figure 5). In other words, a practitioner pays less by overshooting than by undershooting the choice of n_0. The boundary mechanism explains the asymmetry: too few random samples can expose the worst-case boundary itinerary, while too many merely waste a few extra uninformed samples without distorting the overall BO trajectory.

Multi-step lookahead BO as a future direction. The boundary pathology is myopic in nature: a single-step acquisition function has no way to account for the wasted future iterations incurred by a corner query. Multi-step lookahead BO may implicitly account for this cost. Two complementary methods are immediate candidates: Cheon et al. [2025]'s EARL-BO uses reinforcement learning for high-dimensional multi-step BO, and Wu and Frazier [2019] provide a practical two-step lookahead acquisition. We conjecture that even two-step lookahead can help mitigate the initialization tradeoff by pricing in the "next-query cost," while preserving the advantages of variance-driven acquisitions.

References

Mujin Cheon, Haeun Byun, and Jay H Lee. Non-myopic bayesian optimization using model-free reinforcement learning and its application to optimization in electrochemistry. *Computers & Chemical Engineering*, 184:108624, 2024.

Mujin Cheon, Jay H Lee, Dong-Yeun Koh, and Calvin Tsay. EARL-BO: reinforcement learning for multi-step lookahead, high-dimensional bayesian optimization. In *International Conference on Machine Learning*, pages 10182–10198. PMLR, 2025.

Peter I. Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2023.

Mark E Johnson, Leslie M Moore, and Donald Ylvisaker. Minimax and maximin distance designs. *Journal of Statistical Planning and Inference*, 26(2):131–148, 1990.

Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.

Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.

Jason L Loeppky, Jerome Sacks, and William J Welch. Choosing the sample size of a computer experiment: A practical guide. *Technometrics*, 51(4):366–376, 2009.

ChangYong Oh, Efstratios Gavves, and Max Welling. Bock: Bayesian optimization with cylindrical kernels. In *International conference on machine learning*, pages 3868–3877. PMLR, 2018.

Miles Olson, Elizabeth Santorella, Louis C Tiao, Sait Cakmak, Mia Garrard, Samuel Daulton, Zhiyuan Jerry Lin, Sebastian Ament, Bernard Beckerman, Eric Onofrey, et al. Ax: A platform for adaptive experimentation. In *AutoML 2025 ABCD Track*, 2025.

Joel A Paulson and Calvin Tsay. Bayesian optimization as a flexible and efficient design framework for sustainable process systems. *Current Opinion in Green and Sustainable Chemistry*, 51:100983, 2025.

Eero Siivola, Aki Vehtari, Jarno Vanhatalo, Javier González, and Michael Riis Andersen. Correcting boundary over-exploration deficiencies in bayesian optimization with virtual derivative sign observations. *arXiv preprint arXiv:1704.00963*, 2017.

Kevin Swersky. *Improving Bayesian optimization for machine learning using expert priors*. University of Toronto (Canada), 2017.

William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

Jian Wu and Peter I. Frazier. Practical two-step lookahead Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2019.

Yilin Xie, Shiqiang Zhang, Jixiang Qing, Ruth Misener, and Calvin Tsay. Bogrape: Bayesian optimization over graphs with shortest-path encoded. *arXiv preprint arXiv:2503.05642*, 2025.

A Additional Results

A.1 Experimental Details

This subsection expands the brief setup of Section 2.

(E1) Oracle GP, 3D discrete, EI. A single $f \sim \mathrm{GP}(0, k_{5/2})$ with lengthscale $\ell{=}1$ and signal variance $\sigma_f^2{=}1$ is drawn once on $\{-5, -4, \ldots, 4, 5\}^3$ ($11^3 = 1331$ grid points). The BO surrogate uses the *exact* generative kernel and hyperparameters (`optimizer = None`, `length_scale_bounds = fixed`); τ is the discrete first-hit time $\min\{t : x_t = x^\star\}$. We run 10^5 episodes per $n_0 \in \{1, \ldots, 35\}$, with episode randomness coming only from the random initial subset.

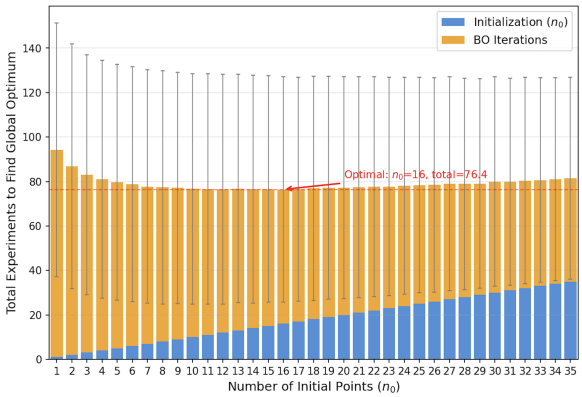

Figure 5: **(E1) Oracle 3D EI.** Total cost averaged over 10^5 episodes with an exact-kernel oracle GP. $n_0^\star$=16, $C(n_0^\star)\approx$76.4; $C(1)\approx$94, $C(35)\approx$81.

(E2) Ackley 4D, RBF, MLE & MCMC, EI / UCB / PI / TS. $f = -\mathrm{Ackley}_4$ on $[-6, 6]^4$ discretised with grid spacing 1.5 ($9^4 = 6561$ grid points). The surrogate uses an isotropic RBF kernel; hyper-parameters are learnt either by L-BFGS MLE with 7 restarts, or by fully Bayesian NUTS-MCMC (warmup 256, samples 64, target acceptance 0.9, max tree depth 10). y values are standardised before fitting and unstandardised at prediction. 5,000 episodes per $n_0 \in \{1, \ldots, 20\}$.

(E3) Oracle 3D grid, four acquisitions. Same domain, kernel, and surrogate as (E1)—fixed $f \sim \mathrm{GP}(0, k_{5/2})$ on $\{-5, \ldots, 5\}^3$ with the exact generative kernel—but extended to all four of EI, UCB (κ=2), PI, and TS, with 500 episodes per $(n_0, \text{acquisition})$ pair and $n_0 \in \{0, 10, 20, 30, 40, 50\}$. Performance is reported as simple regret $f^\star - \max f_{\mathrm{obs}}$ rather than first-hit time so that all four acquisitions can be compared on the same metric. The four acquisitions are evaluated on the *same* per-episode random initializations to remove cross-acquisition variance.

A.2 (E3) Alternative Grouping

Figure 3 of the main text groups the (E3) data by acquisition function (one panel per acquisition, one curve per n_0). Figure 6 below regroups the same data per n_0 so that the four acquisitions can be compared side by side. EI / UCB / PI cluster below TS at every n_0.

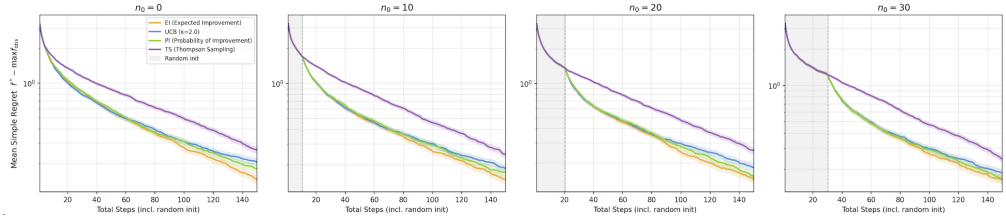

Figure 6: **(E3) Oracle 3D grid, alternative grouping.** Same data as Figure 3, regrouped per n_0 so that the four acquisitions can be compared side by side at every n_0.

A.3 Continuous 3D Corroboration

To corroborate the (E3) regret-based dichotomy outside a fixed discrete grid, we run a continuous 3D experiment with EI, UCB (κ=2), PI, and TS on per-episode random Fourier feature draws of f on $[-5, 5]^3$ (candidate set re-sampled every BO step), 500 episodes per $(n_0, \text{acquisition})$ pair, $n_0 \in \{0, 10, 20, 30, 40\}$, with matched random initializations. Figures 7–8 show the same pattern as (E3): EI / UCB / PI show a lower absolute simple regret with n_0 sensitivity, while Thompson Sampling is essentially n_0-flat at a higher absolute regret level.

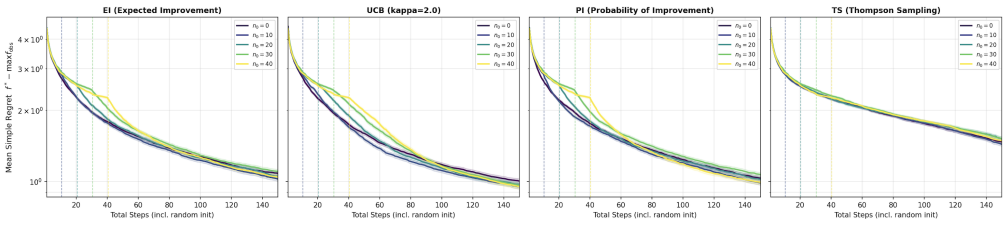

Figure 7: **Continuous 3D Oracle, four acquisitions (panel-per-acquisition view).** Simple regret vs total steps on $[-5, 5]^3$ with per-episode random Fourier feature draws, 500 episodes per $(n_0, \text{acquisition})$ pair, $n_0 \in \{0, 10, 20, 30, 40\}$. EI / UCB / PI show a lower absolute regret with n_0 sensitivity; TS is essentially n_0-flat at a higher absolute regret.

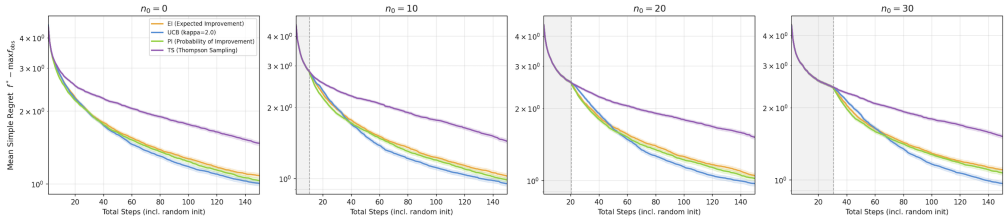

Figure 8: **Continuous 3D Oracle, alternative grouping.** Same data as Figure 7, regrouped per n_0 to compare the four acquisitions side by side.