REIGN: Robust Expected Information Gain for Navigating Adaptive Perturbation Screens

Noel Thomas
Mohamed bin Zayed University of Artificial Intelligence
noel.thomas@mbzuai.ac.ae

Abstract

Perturbation screens test hundreds of compounds or guide RNAs across multiple experimental contexts, but each query costs reagents, instrument time, and analyst attention. Adaptive experiment selection can speed up hit discovery, yet standard Bayesian optimization treats each context in isolation, ignoring shared structure. We propose REIGN, which combines a conjugate hierarchical Gaussian surrogate with a cross-context information-gain acquisition function. The surrogate propagates posteriors from completed contexts to new ones via exponential moving average updates; the acquisition inflates within-context posterior variance by the estimated between-context variability, directing exploration toward actions whose effects are least stable across prior contexts. We evaluate on three cached-oracle benchmarks: SciPlex3 (99 compounds, 3 cell lines), Shifrut2018 (120 CRISPR targets, 4 donor×stimulation conditions), and Buchwald-Hartwig (264 reaction conditions, 15 substrate contexts). On Buchwald-Hartwig, REIGN reaches Hit@1 = 0.351 (95% CI $[0.325, 0.375]$), tied with BO-UCB (0.352) and well above greedy transfer (0.252, $p < 0.001$, n=50 seeds), Thompson sampling (0.303), and max-value entropy search (0.304). Turning off cross-context transfer drops every method to random-level performance (Hit@1 = 0.197, Cohen's d=1.0), confirming that the hierarchical surrogate, not the acquisition rule, drives the gains. On SciPlex3 and Shifrut2018 ($3-4$ contexts), all adaptive methods cluster together and greedy remains competitive, consistent with insufficient contexts for the hierarchical prior to mature. A controlled context-count sweep on Buchwald-Hartwig pins down this transition: exploration-based methods pull ahead of greedy only once roughly 15 contexts have been observed.

1 Introduction

Perturbation screens are central to modern biology and chemistry. They measure the effect of genetic knockouts, small-molecule treatments, or reaction conditions on phenotypic readouts (Shahriari et al., 2015). Screens can test hundreds to thousands of conditions, but cost and throughput constraints make adaptive selection valuable: choosing the next experiment based on what has already been observed can compress the total number of queries needed to find high-value hits.

Most adaptive methods treat each experimental context (cell line, patient sample, substrate) as a fresh problem. This is wasteful when contexts share structure. The best catalytic conditions for one aryl halide substrate often transfer to related substrates; the same gene target may respond similarly across T-cell donors. Hierarchical modelling can exploit this shared structure to accelerate discovery in each new context.

Our central empirical finding is specific: transfer-aware exploration helps only when enough contexts accumulate for the hierarchical prior to become informative. On Buchwald-Hartwig (15 substrate contexts, 264 reaction conditions), exploration-based methods (REIGN and BO-UCB) reach Hit@1 ≈ 0.35 while greedy transfer achieves 0.25 ($p < 0.001$). On SciPlex3 (3 contexts) and Shifrut2018 (4 contexts), all adaptive methods are statistically indistin-

guishable. A controlled sweep over context counts confirms a sharp transition: exploration pulls ahead of greedy only at 15 contexts.

We make three contributions:

1. A conjugate hierarchical Gaussian surrogate that propagates posteriors across streaming contexts via Bayesian updating. Because all updates are in closed form, selection costs $O(|\mathcal{A}|)$ per query with no matrix inversion.
2. A cross-context information-gain acquisition function (REIG) that augments within-context posterior variance with the estimated between-context variability of each action, preferentially exploring actions whose effects are unstable across prior contexts.
3. An evaluation on three cached-oracle benchmarks with six baselines, ablations over all key hyperparameters, noise-injection tests, and a controlled context-count sweep. All code and oracles are released at https://github.com/11NOel11/REIGN.

2 Background

Bayesian optimization. Bayesian optimization (BO) (Mockus, 1978; Jones et al., 1998) fits a probabilistic surrogate to the objective $f : \mathcal{A} \to \mathbb{R}$ and selects queries by maximizing an acquisition function. Expected Improvement (EI) (Mockus, 1978) and Upper Confidence Bound (UCB) (Auer, 2002) are the most common choices. Standard BO restarts from scratch in each new experimental context.

Expected information gain. The expected information gain (EIG) of observing outcome y about a latent parameter θ is (Lindley, 1956)

$$\mathrm{EIG}(\theta; y) = \mathbb{E}_y\big[H(\theta) - H(\theta \mid y)\big], \tag{1}$$

where H is differential entropy. Under a conjugate Gaussian model with posterior variance v and noise variance σ^2, this reduces to

$$\mathrm{EIG}(a; c) = \tfrac{1}{2}\log\big(1 + v_{c,a}/\sigma^2\big). \tag{2}$$

3 Method

3.1 Problem setup

We consider a multi-context screen with contexts $\mathcal{C}$ (substrates, cell lines) and actions $\mathcal{A}$ (reaction conditions, compounds). Contexts arrive in a stream $c_1, c_2, \ldots$; within each context we have a budget of B queries. Each query (c, a) returns a scalar outcome $y_{c,a}$.

3.2 Hierarchical Gaussian surrogate

We place a two-level normal hierarchy over outcomes:

$$\mu_0 \sim \mathcal{N}(0,\, \tau^2\mathbf{I}), \tag{3}$$

$$\mu_{c,a} \sim \mathcal{N}(\mu_{0,a},\, \eta^2), \tag{4}$$

$$y_{c,a} \mid \mu_{c,a} \sim \mathcal{N}(\mu_{c,a},\, \sigma^2). \tag{5}$$

After observing $n_{c,a}$ measurements with sample mean $\bar{y}_{c,a}$ in context c, the posterior over $\mu_{c,a}$ is

$$\mu_{c,a} \mid \mathcal{D}_c \;\sim\; \mathcal{N}(m_{c,a},\, v_{c,a}), \tag{6}$$

where $m_{c,a}$ and $v_{c,a}$ follow from standard normal-normal conjugacy.

Posterior-as-prior propagation. When context c ends, we update the global prior mean for each queried action using an exponential moving average (EMA) with weight α:

$$\hat{\mu}_{0,a} \;\leftarrow\; (1-\alpha)\,\hat{\mu}_{0,a} \;+\; \alpha\, m_{c,a}. \tag{7}$$

The between-context variance s_a^2 is tracked simultaneously via Welford's online algorithm over the sequence of posterior means $m_{1,a}, m_{2,a}, \ldots$ Only queried actions are updated; unqueried actions retain their prior, avoiding artificial variance collapse.

3.3 Cross-context information gain (REIG)

Standard EIG (Eq. 2) uses only the within-context posterior variance $v_{c,a}$. When an action's effect varies across prior contexts, this understates true uncertainty: the transferred prior mean may be unreliable precisely where cross-context variability is large. We augment the posterior variance with the Welford-estimated between-context variance:

$$v_{c,a}^{\text{aug}} \;=\; v_{c,a} \;+\; \rho\, s_a^2, \tag{8}$$

where $\rho \geq 0$ controls how much cross-context spread inflates the effective uncertainty. The resulting acquisition function is

$$\text{REIG}(a; c) \;=\; \tfrac{1}{2} \log\!\big(1 + v_{c,a}^{\text{aug}}/\sigma^2\big). \tag{9}$$

Setting $\rho = 0$ recovers plain EIG. For $\rho > 0$, actions with high s_a^2 receive an exploration bonus. On Buchwald-Hartwig, cross-context variability correlates positively with mean yield ($r = 0.36$), so REIG preferentially explores high-yield, substrate-specific conditions.

3.4 REIGN planner

The full REIGN score combines information gain with exploitation:

$$\text{score}(a; c) \;=\; \lambda \cdot \widetilde{\text{REIG}}(a; c) \;+\; (1-\lambda) \cdot \widetilde{\text{EI}}(a; c), \tag{10}$$

where tildes denote min-max normalization over the candidate set and $\lambda \in [0, 1]$ controls the balance. We fix $\lambda = 0.5$ throughout.

4 Experiments

4.1 Datasets

SciPlex3. A chemical perturbation screen (Srivatsan et al., 2020): 99 compounds across 3 cell lines (A549, K562, MCF7). Target: z-scored 10-gene proliferation signature. Budget 30 per context, warm-start 3, averaged over 30 seeds.

Shifrut2018. A CRISPR activation screen in primary human T cells (Shifrut et al., 2018): 120 gene targets across 4 donor×stimulation contexts. Target: 10-gene T-cell activation signature. Budget 30, 30 seeds.

Buchwald-Hartwig. A reaction optimization benchmark (Ahneman et al., 2018): 264 conditions (4 ligands × 3 bases × 22 additives) across 15 aryl halide substrates. Target: yield normalized to $[0, 1]$. Budget 50, warm-start 3, 50 seeds. This is the primary testbed because it has the most contexts.

Baselines. Random (uniform); Greedy (highest posterior mean, with EMA transfer); BO-EI and BO-UCB (standard BO); Thompson sampling (Russo & Van Roy, 2016); max-value entropy search (MES) (Wang & Jegelka, 2017). All methods share the same hierarchical surrogate and transfer mechanism (α=0.9, σ^2=0.1, τ_0^2=1.0). Only the acquisition function differs.

Table 1: Buchwald-Hartwig (n=50 seeds, budget=50, 15 substrates). REIGN and BO-UCB are tied on Hit@1 and both strongly outperform Greedy. Bold: best per column.

Method	Hit@1	95% CI	Hit@3	Hit@5	nDCG@1	AUC
Random	0.197	[0.165, 0.231]	0.463	0.665	0.938	0.104
Greedy	$0.252^{\dagger}$	[0.215, 0.289]	0.571	0.759	0.956	0.195
BO-EI	0.327^{**}	[0.297, 0.356]	0.615	0.757	0.948	0.197
BO-UCB	0.352	[0.328, 0.376]	0.617	0.749	0.947	0.204
REIGN	0.351	[0.325, 0.375]	0.620	0.751	0.947	0.204
REIGN (no transfer)	0.197	[0.165, 0.231]	0.463	0.665	0.938	0.104

Superscripts: REIGN vs. each baseline (paired t-test, $n = 50$ seeds): $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$. Greedy: $p < 0.001$, $d = 0.58$. BO-EI: $p = 0.001$, $d = 0.48$. BO-UCB: $p = 0.66$ (tied). No-transfer: $p < 0.001$, Cohen's $d = 1.0$.

Table 2: SciPlex3 (n=30 seeds, 3 cell-line contexts). All adaptive methods converge to the same Hit@1. Bold: best per column.

Method	Hit@1	95% CI	AUC
Random	0.300	[0.211, 0.400]	0.154
Greedy	0.456	[0.356, 0.567]	0.272
BO-EI	0.456	[0.356, 0.567]	0.275
BO-UCB	0.456	[0.356, 0.567]	0.275
REIGN	0.456	[0.356, 0.567]	0.275

4.2 Results

Buchwald-Hartwig (15 contexts). Table 1 shows the main results. REIGN (Hit@1=0.351) and BO-UCB (0.352) are statistically tied (p=0.66) and both significantly outperform Greedy (0.252, p<0.001, Cohen's d=0.58), BO-EI (0.327, p=0.001), Thompson (0.303), and MES (0.304). See Appendix A for the full baseline comparison.

Transfer is the main driver. When we disable EMA propagation (each context starts from a flat prior), Greedy, BO-UCB, and REIGN all collapse to random-level performance (Hit@1=0.197, p<0.001 vs. REIGN with transfer, Cohen's d=1.0). The acquisition function matters secondarily: once transfer is on, REIG and UCB produce nearly identical results. The hierarchical surrogate is doing the heavy lifting.

When does exploration help? Figure 2 shows per-context Hit@1 trajectories. Through contexts 1 – 10, Greedy, REIGN, and BO-UCB perform comparably. Starting at context 11, exploration-based methods jump to Hit@1 $\approx$ 0.70 while Greedy stays near 0.20. This suggests that the hierarchical prior needs roughly 10 contexts before it becomes informative enough for exploration to exploit.

We test this directly with a controlled context-count sweep (Figure 1). With 1 – 10 contexts, Greedy matches or beats REIGN and BO-UCB. At 15 contexts, exploration-based methods pull sharply ahead. This pins down the transition: transfer-aware exploration becomes useful only after the prior has absorbed enough cross-context evidence.

SciPlex3 and Shifrut2018 (3 – 4 contexts). Tables 2 and 3 confirm the expected boundary behavior. On SciPlex3, every adaptive method reaches Hit@1=0.456, identical within confidence intervals. On Shifrut2018, Greedy and BO both achieve 0.325 while REIGN reaches 0.292; with only 4 contexts, the exploration overhead is counterproductive. These are not failure cases for REIGN but rather the regime where the hierarchical prior has too few contexts to mature.

Table 3: Shifrut2018 (n=30 seeds, 4 donor×stimulation contexts). With few contexts, greedy exploitation is competitive.

Method	Hit@1	95% CI	AUC
Random	0.133	[0.083, 0.192]	0.070
Greedy	0.325	[0.225, 0.433]	0.227
BO-EI	0.325	[0.225, 0.433]	0.230
BO-UCB	0.325	[0.225, 0.433]	0.230
REIGN	0.292	[0.183, 0.400]	0.206

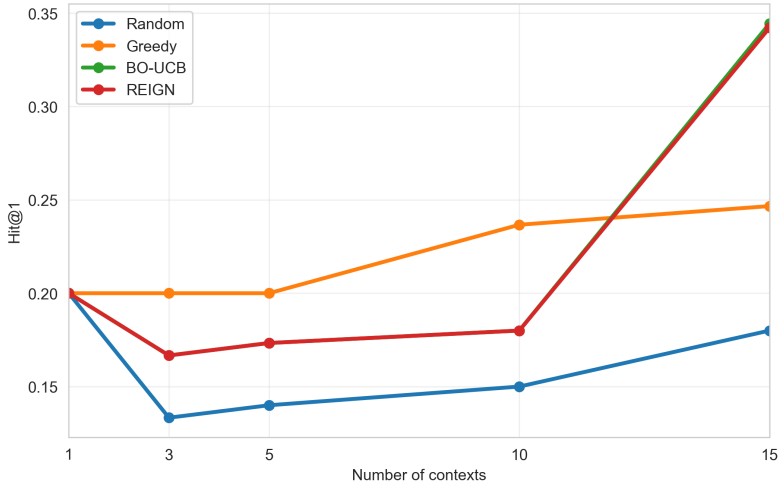

Figure 1: Context-count sweep on Buchwald-Hartwig (n=30 seeds). Greedy matches or beats exploration-based methods through 10 contexts. At 15 contexts, REIGN and BO-UCB pull ahead, consistent with the hierarchical prior needing enough observations to mature.

5 Discussion

Exploration vs. exploitation. Our results give a concrete answer to when exploration-based acquisition outperforms greedy exploitation: on Buchwald-Hartwig (15 contexts, 264 actions), the difference is large and significant (p<0.001, d=0.58). On SciPlex3 (3 contexts), there is no difference. On Shifrut2018 (4 contexts), Greedy actually wins (p=0.043). The context-count sweep (Figure 1) confirms this is a regime effect, not noise.

REIGN vs. BO-UCB. On Buchwald-Hartwig, REIGN and BO-UCB produce nearly identical results. The cross-context variance augmentation ($\rho > 0$ in Eq. 8) provides a small, consistent benefit over plain EIG (Hit@1: 0.351 vs. 0.347, Appendix D), but confidence intervals overlap. We do not claim that REIG is clearly superior to UCB; the shared surrogate and transfer mechanism are what matter most. REIG offers a principled information-theoretic motivation for cross-context exploration, but the practical gap is small.

Deterministic oracles. All three benchmarks use pre-computed, noiseless outcomes. Appendix B reports results under synthetic noise injection. REIGN and BO-UCB remain tied at moderate noise (σ=0.05) and both degrade to random at high noise (σ=0.2). Paired tests are non-significant at all noise levels (p>0.4). Testing on data with genuine biological replicate variance is important future work.

Search space scale. Our benchmarks contain 99 – 264 actions. Real screens test thousands to millions of conditions. REIGN's $O(|\mathcal{A}|)$ per-selection cost (closed-form posteriors, no matrix inversion) scales in principle, but validating this requires larger benchmarks.

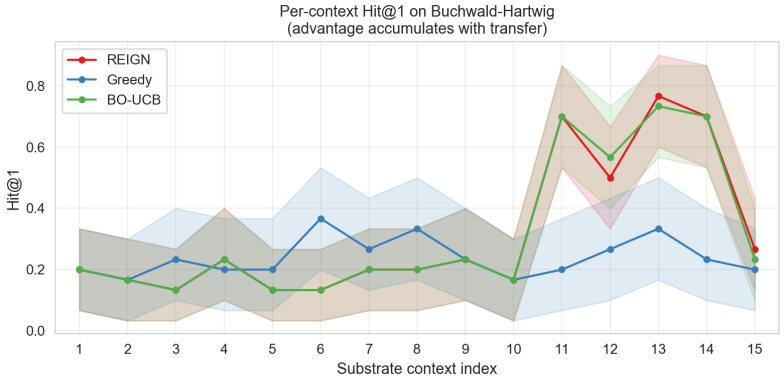

Figure 2: Per-context Hit@1 on Buchwald-Hartwig (n=30 seeds). Through contexts 1–10, methods are comparable. From context 11 onward, exploration-based methods reach Hit@1 ≈ 0.70 while Greedy stays near 0.20. Bands: 95% bootstrap CIs.

Action structure. REIGN treats actions as exchangeable. It transfers information across contexts but ignores similarity between actions (e.g., structural similarity between compounds). Adding action-level kernels or learned embeddings to the surrogate is a natural extension.

Limitations. We do not evaluate on wet-lab experiments. The benchmarks cover a narrow slice of screening tasks. The surrogate assumes Gaussian outcomes and independent actions, which may not hold in practice. REIGN requires enough prior contexts to estimate meaningful between-context variability; with fewer than roughly 10 contexts, greedy exploitation is preferable.

6 Conclusion

REIGN combines a conjugate hierarchical Gaussian surrogate with a cross-context information-gain acquisition function for multi-context perturbation screens. On Buchwald-Hartwig (15 substrates, 264 conditions, 50 seeds), it matches the strongest BO baseline while outperforming greedy transfer, Thompson sampling, and MES. Disabling cross-context transfer collapses all methods to random, showing that the hierarchical surrogate is the main source of improvement. On biological screens with 3–4 contexts, simpler methods remain competitive, which marks the boundary of the useful regime rather than a hidden failure case. Code, cached oracles, and experiment configurations are at https://github.com/11NOel11/REIGN.

References

Derek T Ahneman, Jesús G Estrada, Shishi Lin, Spencer D Dreher, and Abigail G Doyle. Predicting reaction performance in C–N cross-coupling using machine learning. Science, 360(6385):186–190, 2018.

Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. Journal of Machine Learning Research, 3:397–422, 2002.

Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. Journal of Global Optimization, 13(4):455–492, 1998.

Dennis V Lindley. On a measure of the information provided by an experiment. The Annals of Mathematical Statistics, 27(4):986–1005, 1956.

Jonas Mockus. Toward Bayesian optimal design of experiments. Towards Global Optimization, 2, 1978.

Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of Thompson sampling. Journal of Machine Learning Research, 17(1):2442–2471, 2016.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization. Proceedings of the IEEE, 104(1):148–175, 2015.

Eric Shifrut, Julia Carnevale, Victoria Tobin, Theodore L Roth, Jonathan M Woo, Catherine T Bui, P Jonathan Li, Morgan E Diolaiti, Alan Ashworth, and Alexander Marson. Genome-wide CRISPR screens in primary human T cells reveal key regulators of immune function. Cell, 175(7):1958–1971, 2018.

Sanjay R Srivatsan, José L McFaline-Figueroa, Vijay Ramani, Lauren Saunders, Junyue Cao, Jonathan Packer, Hannah A Pliner, Dana L Jackson, Riza M Daza, Lena Christiansen, et al. Massively multiplex chemical transcriptomics at single-cell resolution. Science, 367(6473):45–51, 2020.

Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient Bayesian optimization. In International Conference on Machine Learning, pp. 3627–3635, 2017.

A Transfer Attribution and Additional Baselines

Table 4 shows the full transfer-on / transfer-off comparison on Buchwald-Hartwig. With transfer disabled, Greedy, BO-UCB, and REIGN all produce results numerically identical to Random on every seed. Thompson sampling and MES are included as additional exploration baselines; both fall below REIGN and BO-UCB.

Table 4: Transfer attribution on Buchwald-Hartwig (n=50 seeds). "Transfer off" disables EMA propagation between contexts. All transfer-off results match Random exactly.

Method	Transfer on	Transfer off	AUC (on)
Greedy	0.252	0.197	0.195
BO-UCB	0.352	0.197	0.204
REIGN	0.351	0.197	0.204
Thompson	0.303	–	0.157
MES	0.304	–	0.178
Random	0.197	0.197	0.104

B Noise-Injection Experiments

We add Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ to oracle responses while keeping ground-truth rankings noiseless for metric evaluation.

Table 5: Hit@1 under oracle noise (n=50 seeds, Buchwald-Hartwig). Only the observations seen by the planner are corrupted.

Method	σ=0.0	σ=0.05	σ=0.1	σ=0.2
Random	0.197	0.113	0.069	0.037
Greedy	0.252	0.129	0.056	0.032
BO-UCB	0.352	0.131	0.083	0.037
REIGN	0.351	0.133	0.076	0.036

REIGN and BO-UCB are tied at σ=0.0. At σ=0.05, REIGN is marginally ahead; at σ=0.1, BO-UCB is marginally ahead. By σ=0.2 every method is near random. Paired t-tests between REIGN and BO-UCB are non-significant at all noise levels (all $p > 0.4$). We interpret the noise robustness of both methods as modest and roughly equivalent.

C Top-k Batch Comparison

Table 6 compares batch variants that select k actions per round before updating the surrogate.

Table 6: Top-k on Buchwald-Hartwig (n=30 seeds, budget=50). k actions are queried simultaneously per step.

k	Method	Hit@1	nDCG@1	AUC
1	UCB	0.344	0.950	0.195
1	REIGN	0.353	0.946	0.205
3	UCB	0.351	0.945	0.195
3	REIGN	0.353	0.945	0.192
5	UCB	0.353	0.948	0.211
5	REIGN	0.353	0.948	0.207

REIGN has a slight edge at k=1 and k=3; by k=5 the two methods converge, consistent with the batch setting becoming saturated.

D Parameter Ablation

Table 7 reports sensitivity to λ (REIG vs. EI weight) and ρ (cross-context augmentation strength). Table 8 varies the EMA transfer weight α. All runs use n=50 seeds on Buchwald-Hartwig.

Table 7: λ and ρ ablation (n=50 seeds). Top: λ sweep with ρ=1.0. Bottom: ρ sweep with λ=0.5. Bold: best.

Config	Hit@1	95% CI
$\lambda = 0.00$, $\rho = 1.0$ (pure EI)	0.327	[0.297, 0.356]
$\lambda = 0.25$, $\rho = 1.0$	0.328	[0.301, 0.355]
$\lambda = 0.50$, $\rho = 1.0$ (default)	0.351	[0.327, 0.375]
$\lambda = 0.75$, $\rho = 1.0$	0.304	[0.280, 0.329]
$\lambda = 1.00$, $\rho = 1.0$ (pure REIG)	0.273	[0.240, 0.307]
$\lambda = 0.50$, $\rho = 0.0$ (plain EIG)	0.347	[0.323, 0.371]
$\lambda = 0.50$, $\rho = 0.1$	0.349	[0.325, 0.375]
$\lambda = 0.50$, $\rho = 0.5$	0.351	[0.327, 0.375]
$\lambda = 0.50$, $\rho = 1.0$ (default)	0.351	[0.327, 0.375]
$\lambda = 0.50$, $\rho = 2.0$	0.349	[0.325, 0.373]

Table 8: EMA weight α ablation (n=50 seeds). Higher α gives more weight to the most recent context posterior.

EMA weight α	Hit@1	95% CI
0.1	0.321	[0.293, 0.351]
0.3	0.328	[0.301, 0.356]
0.5	0.329	[0.304, 0.355]
0.7	0.344	[0.319, 0.369]
0.9	0.351	[0.327, 0.375]

Three observations. (1) λ sensitivity is non-monotone: pure EI (λ=0, Hit@1=0.327) and pure REIG (λ=1, 0.273) both fall below the balanced combination (λ=0.5, 0.351). Information-gain exploration and greedy exploitation are complementary. (2) Plain EIG (ρ=0) already reaches 0.347; positive ρ values cluster around 0.349 – 0.351 with overlapping confidence intervals. The cross-context bonus helps, but the effect is secondary to transfer. (3) Hit@1 increases monotonically with α, from 0.321 at α=0.1 to 0.351 at α=0.9. Aggressive posterior-as-prior propagation is better on this benchmark, consistent with the 15 substrates being related enough that recent context posteriors are informative priors for the next.