

---

# PriorBand: Practical Hyperparameter Optimization in the Age of Deep Learning

---

**Neeratoy Mallik**  
University of Freiburg  
mallik@cs.uni-freiburg.de

**Edward Bergman**  
University of Freiburg  
bergmane@cs.uni-freiburg.de

**Carl Hvarfner**  
Lund University  
carl.hvarfner@cs.lth.se

**Danny Stoll**  
University of Freiburg  
stolld@cs.uni-freiburg.de

**Maciej Janowski**  
University of Freiburg  
janowski@cs.uni-freiburg.de

**Marius Lindauer**  
Leibniz University Hannover  
m.lindauer@ai.uni-hannover.de

**Luigi Nardi**  
Lund University  
Stanford University  
DBtune  
luigi.nardi@cs.lth.se

**Frank Hutter**  
University of Freiburg  
fh@cs.uni-freiburg.de

## Abstract

Hyperparameters of Deep Learning (DL) pipelines are crucial for their downstream performance. While a large number of methods for Hyperparameter Optimization (HPO) have been developed, their incurred costs are often untenable for modern DL. Consequently, manual experimentation is still the most prevalent approach to optimize hyperparameters, relying on the researcher’s intuition, domain knowledge, and cheap preliminary explorations. To resolve this misalignment between HPO algorithms and DL researchers, we propose **PriorBand**, an HPO algorithm tailored to DL, able to utilize both expert beliefs and cheap proxy tasks. Empirically, we demonstrate **PriorBand**’s efficiency across a range of DL benchmarks and show its gains under informative expert input and robustness against poor expert beliefs.

## 1 Introduction

The performance of Deep Learning (DL) models crucially depends on dataset-specific settings of their hyperparameters (HPs) [1, 2]. Therefore, Hyperparameter Optimization (HPO) is an integral step in the development of DL models [3–6]. HPO methods have typically been applied to traditional machine learning models (including shallow neural networks) that focus on fairly small datasets [7, 8]. Current DL practitioners, however, often utilize much larger models and datasets (e.g., a single training of GPT-3 [9] was estimated to require  $3 \cdot 10^{23}$  FLOPS, i.e., months on a thousand V100 GPUs [10]). As recent HPO practices continue to apply simple techniques like grid, random, or manual search [11–14], existing HPO methods seem misaligned with DL practice.

To make the misalignment between existing HPO methods and DL practice explicit, we identify the following desiderata for an efficient, scalable HPO method suitable for current DL:

Table 1: Comparison with respect to identified desiderata for DL pipelines. Shown is our algorithm, PriorBand, along with Grid Search, multi-fidelity random search (MF-RS), e.g. HyperBand, ASHA, Bayesian Optimization (BO), BO with Expert Priors ( $\pi$ BO), multi-fidelity BO (MF-BO), e.g. BOHB, Mobster. A checkmark and red cross indicates whether they satisfy the corresponding desideratum. Parenthesized checkmark indicates the desideratum is partly fulfilled, or requires additional modifications to be fulfilled. Model-based methods require customized kernels for discrete hyperparameters and asynchronous parallel approaches [7, 15] to achieve speed-up under (asynchronous) parallelism. Multi-fidelity algorithms’ final performance is contingent on the budget spent on HPO.

Criterion	Grid Search	MF-RS	BO	BO with Priors	MF-BO	PriorBand
Good anytime performance	✗	✓	✗	✗	✓	✓
Good final performance	✗	(✓)	✓	✓	✓	✓
Proxy task integration	✗	✓	✗	✗	✓	✓
Integrate expert beliefs	✗	✗	✗	✓	✗	✓
Mixed search spaces	✓	✓	(✓)	(✓)	(✓)	✓
Model free	✓	✓	✗	✗	✗	✓
Speedup under parallelism	✓	✓	(✓)	(✓)	(✓)	✓

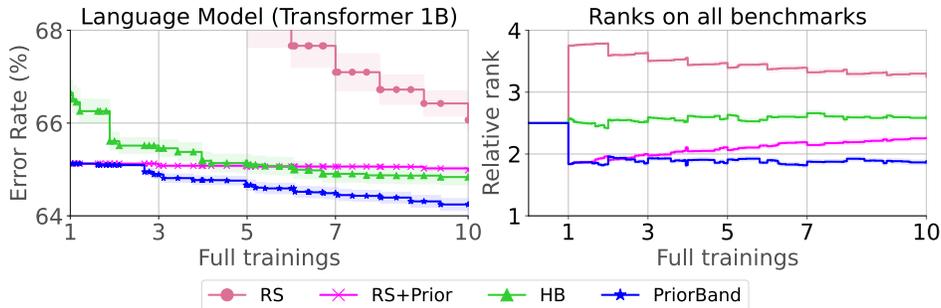


Figure 1: Both plots compare Random Search (RS), sampling from prior (RS+Prior), HyperBand (HB), and our method PriorBand which utilizes a *good* prior as defined by an expert. **[Left]** Tuning a large transformer on the 1B word benchmark. PriorBand leverages the prior, achieving strong anytime performance. **[Right]** Ranks on all our 12 benchmarks; PriorBand consistently ranks best.

1. **Strong performance under low compute budgets:** As DL pipelines become larger and more complex, only a few model training can be afforded to find performant configurations.
2. **Integrate cheap proxy tasks:** To save resources, cheap evaluations must be used to find promising configurations early while maintaining robustness to uninformative proxies.
3. **Integrate expert beliefs:** DL models often come with a competitive default setting or an expert may have prior beliefs on what settings to consider for HPO. Incorporating this information is essential while maintaining robustness to such beliefs being inaccurate.
4. **Handle mixed type search spaces:** Modern DL pipelines tend to be composed of many kinds of hyperparameters, e.g. categorical, numerical, and log hyperparameters.
5. **Simple to understand and implement:** The HPO method should be easy to apply to different DL pipelines and the behavior should be conceptually easy to understand.
6. **Parallelism:** Scaling to modern compute resources must be simple and effective.

Existing HPO algorithms satisfy subsets of these desiderata (see Table 1). For example, methods that utilize cheaper proxy tasks exploit training data subsets [16], fewer epochs/updates [17, 18], and proxy performance predictors [19, 20]. Previous works also consider how to combine expert knowledge into the optimization procedure in the form of the optimal value [21], meta or transfer learned knowledge [22–25] or configurations and regions that are known to the expert to have worked well previously [26–28]. Nonetheless, no HPO method exists that meets all of the desiderata above, especially, a simple model-free approach that can run cheaper proxy tasks, letting the DL expert integrate their domain knowledge for HPO.

We propose PriorBand, the first method to fulfill all desiderata. Our **contributions** are as follows:

1. We are the first to develop an approach to leverage cheap proxy tasks with an expert prior input (Section 2), and we show the need beyond a naive solution to do so (Section 4).
2. We contribute a general HPO algorithm with `PriorBand`, fulfilling all desiderata (Table 1) required for application to DL (Section 5; also see Figure 1)
3. We demonstrate the efficiency and robustness of `PriorBand` on a wide suite of DL tasks under practically feasible compute resources (Section 7).
4. We highlight the flexibility of the method behind `PriorBand` by improving different multi-fidelity algorithms and their model-based extensions (Section 7.2 & 7.3), which highlights the flexible, modular design and possible extensions of `PriorBand`.

We conclude, highlight limitations, and discuss further research directions in Section 10. Our code for reproducing the experiments is open-sourced at <https://github.com/automl/mf-prior-exp>.

## 2 Problem statement

We consider the problem of *minimizing* an expensive-to-evaluate objective function  $f : \Lambda \rightarrow \mathbb{R}$ , i.e.,  $\lambda^* \in \arg \min_{\lambda \in \Lambda} f(\lambda)$ , where a configuration  $\lambda$  comes from a search space  $\Lambda$  which may consist of any combination of continuous, discrete and categorical variables. In HPO particularly,  $f$  refers to training and validating some model based on hyperparameters defined by  $\lambda$ .

Multi-fidelity optimization of  $f(\lambda)$  requires a proxy function, namely  $\hat{f}(\lambda, z)$ , that provides a cheap approximation of  $f$  at a given fidelity  $z$ , e.g., the validation loss of a model only trained for  $z$  epochs. Our methodology considers a fidelity scale  $z \in [z_{\min}, z_{\max}]$  such that evaluating at  $z_{\max}$  coincides with our true objective:  $f(\lambda) = \hat{f}(\lambda, z_{\max})$ .

To take all desiderata into account, we additionally include a user-specified belief  $\pi : \Lambda \rightarrow \mathbb{R}$ , where  $\pi(\lambda) = \mathbb{P}(\lambda = \lambda^*)$  represents a subjective probability that a given configuration  $\lambda$  is optimal. Thus, the overall objective is

$$\lambda^* \in \arg \min_{\lambda \in \Lambda} \hat{f}(\lambda, z_{\max}), \quad \text{guided by } \pi(\lambda). \quad (1)$$

Our problem is to efficiently solve Equation 1 under a constrained budget, e.g., 10 full trainings of a DL model. Given DL training can diverge, the *incumbent* we return is the configuration-fidelity pair  $(\lambda, z)$  that achieved the lowest error while optimizing for Equation 1.

## 3 Background

While we are the first to target Equation 1 in its full form, below, we introduce the necessary background on methods that use either multi-fidelity optimization *or* expert beliefs.

**Successive Halving (SH)** optimizes Equation 1 (sans  $\pi$ ) as a best-arm identification problem in a multi-armed bandit setting [29]. Given a lower and upper fidelity bound ( $[z_{\min}, z_{\max}]$ ) and a reduction factor ( $\eta$ ), SH discretizes the fidelity range geometrically with a log-factor of  $\eta$ . For example, if the inputs to SH are  $z \in [3, 81]$  and  $\eta = 3$ , the fidelity levels in SH are  $z = \{3, 9, 27, 81\}$  with accompanying enumeration called rungs,  $r = \{0, 1, 2, 3\}$ . Any  $z < z_{\max}$  is called a lower fidelity, a cheaper proxy task in DL training. Every iteration of each round of SH involves uniform random sampling of  $\eta^{s_{\max}-1}$  configurations at the lowest fidelity or  $r = 0$ , where  $s_{\max} = \lfloor \log_{\eta}(z_{\max}/z_{\min}) \rfloor$ . After evaluating all samples, only the top-performing  $(1/\eta)$  configurations are retained, discarding the rest. Subsequently, these  $\eta^{s_{\max}-1}/\eta$  configurations are evaluated at  $r = 1$ , or the next higher discretized fidelity. This is repeated till there is only one surviving configuration at the highest  $r$ . Sampling and evaluating many configurations for cheap, performing a tournament selection to retain strong samples, and repeating this, ensures that the high fidelity evaluations happen for relatively stronger configurations only. SH proves that such an early-stopping strategy guarantees a likely improvement over random search under the same total evaluation budget. Due to its random sampling component, SH is an example of MF-RS in Table 1. See Appendix D.2 for an illustrative example.

**HyperBand (HB)** attempts to address SH’s susceptibility to poor fidelity performance correlations [30]. It has the same HPs as SH, and thus given SH’s fidelity discretization, HB iteratively executes multiple rounds of SH where  $z_{\min}$  is set to the next higher discrete fidelity. Using the

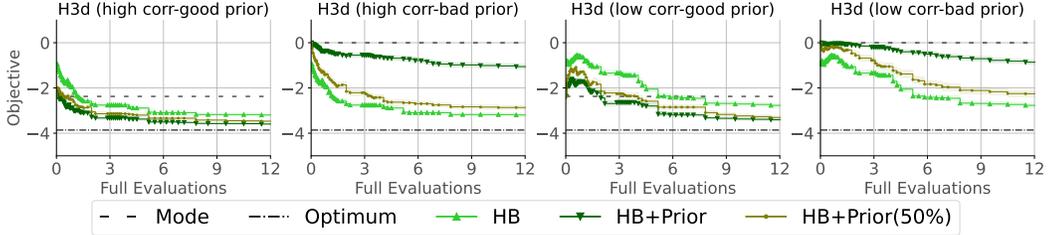


Figure 2: A comparison of naive solutions to Equation 1 on the 3-dimensional multi-fidelity Hartmann benchmarks. We compare different versions of HB, utilizing different strengths of the prior distribution  $\pi(\cdot)$ : HB with 0% influence of the prior (HB), HB with 50% sampling from the prior (HB+Prior(50%)) and HB with 100% sampling from prior (HB+Prior).

example from SH, given  $z \in [3, 81]$  and  $\eta = 3$ , one full iteration of HB corresponds to running  $\text{SH}(z_{\min} = 3, z_{\max} = 81)$  followed by  $\text{SH}(z_{\min} = 9, z_{\max} = 81)$ ,  $\text{SH}(z_{\min} = 27, z_{\max} = 81)$  and  $\text{SH}(z_{\min} = 81, z_{\max} = 81)$ . Notably, such different instantiations of SH in HB do not share any information with each other implicitly. Each execution of  $\text{SH}(z_{\min}, z_{\max})$  is called an SH *bracket*. The sequence of all unique SH brackets run by HB in one iteration is called an HB *bracket*.

$\pi\text{BO}$  utilizes the unnormalized probability distribution  $\pi(\lambda)$  in a BO context to accelerate the optimization using the knowledge provided by the user [28]. In  $\pi\text{BO}$ , the initial design is sampled entirely from  $\pi$  until Bayesian Optimization (BO) begins, where the acquisition function  $\alpha$  is augmented with a prior term that decays over time  $t$ :  $\alpha_{\pi}^t(\lambda) = \alpha(\lambda) \cdot \pi(\lambda)^{\frac{\beta}{t}}$ , where  $\beta$  is an HP. In this work, we borrow  $\pi(\cdot)$  as the expert prior interface to the HPO problem.  $\pi\text{BO}$  relies on the decaying prior on the acquisition  $\alpha_{\pi}^t(\lambda)$  to gradually add more weight to the model and thus allow recovery from a poor initial design under a bad prior. We choose the model-free setting and adapt the strength of the prior based on the evidence of its usefulness.

#### 4 The need to move beyond a naive solution

To solve Equation 1, an intuitive approach is to augment an existing multi-fidelity algorithm, e.g., HyperBand (HB), with sampling from the prior  $\pi(\cdot)$ . In this section, however, we show that this naive approach is not sufficient, motivating the introduction of our algorithm, PriorBand, in Section 5.

To study the naive combination of HyperBand with prior-based sampling, we use the parameterized 3-dimensional multi-fidelity Hartmann benchmarks [31] (Appendix D.1.1) and compare HB with uniform sampling, HB with 100% prior-based sampling, and HB with 50% uniform and prior-based sampling (Figure 2). Unsurprisingly, 100% sampling from priors works best when the prior is helpful and not sampling from the prior works best for misleading priors. 50% random sampling works quite robustly, but is never competitive with the best approach. To achieve robustness to the quality of the prior *and* rival the best approach for each case, we introduce PriorBand below.

#### 5 PriorBand: Moving beyond a naive solution

The key idea behind PriorBand is to complement the sampling strategies used in the naive solution, uniform sampling, and prior-based sampling, with a third strategy: incumbent-based sampling. Thereby, we overcome the robustness issues of the naive solutions presented in Section 4, while still fulfilling the desideratum for simplicity. We first introduce and motivate sampling around the incumbent (Section 5.1), and then describe the ensemble sampling policy  $\mathcal{E}_{\pi}$  that combines all three sampling strategies and incorporates  $\mathcal{E}_{\pi}$  into HB (Section 5.2) for PriorBand.

##### 5.1 Incumbent-based sampling strategy, $\hat{\lambda}(\cdot)$

PriorBand leverages the current incumbent to counter uninformative priors while supporting good priors, as the current incumbent can be seen as indicating a likely good region to sample from. Note that this view on the region around the incumbent is close to the definition of  $\pi(\cdot)$  in  $\pi\text{BO}$  (Section

3), where the prior distribution encodes the expert’s belief about the location of the global optima and thus a good region to sample from.

To construct the incumbent-based sampler  $\hat{\lambda}(\cdot)$ , we perform a local perturbation of the current best configuration. Each hyperparameter (HP) is chosen with probability  $p$  for perturbation. If chosen, continuous HPs are perturbed by  $\epsilon \sim \mathcal{N}(0, \sigma^2)$ . For discrete HPs, we resample with a uniform probability for each categorical value except the incumbent configuration’s value which has a higher probability of selection, discussed further in Appendix E.2.5. For `PriorBand`, we fix these values at  $p = 0.5$  and  $\sigma = 0.25$ . This perturbation operation is simple, easy to implement, and has constant time complexity. We show ablations with two other possible local search designs in Appendix E.2.3.

## 5.2 PriorBand: The $\mathcal{E}_\pi$ -augmented HyperBand

`PriorBand` exploits HB for scheduling and replaces its random sampling component with a combination of random sampling, prior-based sampling, and incumbent-based sampling. We denote the proportions of these individual sampling components as  $p_U$ ,  $p_\pi$ , and  $p_{\hat{\lambda}}$ , respectively, and their combination as the *ensemble sampling policy (ESP)*  $\mathcal{E}_\pi$ .

Figure 3 (left) illustrates `PriorBand` as an extension of HB that, next to the HB hyperparameters ( $z_{\min}$ ,  $z_{\max}$ ,  $\eta$ , `budget`) accepts the expert prior  $\pi$  as an additional input and uses the ESP  $\mathcal{E}_\pi$  in lieu of random sampling. This sampling from  $\mathcal{E}_\pi$  is illustrated in Algorithm 1. Note that  $\mathcal{E}_\pi$  has access to the optimization state ( $s_t$ ) and can thus reactively adapt its individual sampling probabilities  $p_U$ ,  $p_\pi$ , and  $p_{\hat{\lambda}}$  based on the optimization history. We now discuss how we decay the random sampling probability  $p_U$  (Section 5.2.1); and the proportion of incumbent and prior sampling (Section 5.2.2).

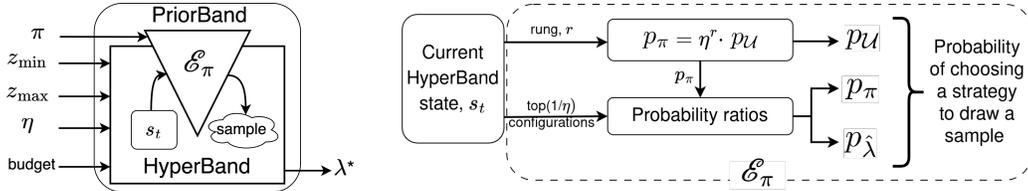


Figure 3: `PriorBand` schema; **[Left]** in the base algorithm vanilla-HB, we replace the random sampling module by  $\mathcal{E}_\pi$ , which can interface an expert prior  $\pi$  and access the state of HB; **[Right]**  $\mathcal{E}_\pi$  reads the state every iteration and determines the probabilities for selecting a sampling strategy.

---

### Algorithm 1 Sampling from $\mathcal{E}_\pi$

---

- 1: **Input:**  $s$ ,  $s_{\max}$ ,  $\eta$ , observations  $\mathcal{H}$ , prior  $\pi$
  - 2:  $r = s_{\max} - s \triangleright s_{\max}, s$  input from HB
  - 3:  $p_U = 1/(1 + \eta^r)$ ,  $p_\pi = 1 - p_U$ ,  $p_{\hat{\lambda}} = 0$
  - 4: **if** `activate_incumbent`( $s_t$ ) **then**
  - 5:      $p_\pi, p_{\hat{\lambda}} \leftarrow$  Algorithm 2( $\mathcal{H}, r, p_\pi$ )
  - 6: **end if**
  - 7:  $d(\cdot) \leftarrow$  sample strategy by  $\{p_U, p_\pi, p_{\hat{\lambda}}\}$
  - 8:  $\lambda \leftarrow$  sample from  $d(\cdot)$
  - 9: **return**  $\lambda$
- 

---

### Algorithm 2 Dynamic weighting of $w_{\hat{\lambda}}$

---

- 1: **Input:** observations  $\mathcal{H}$ , rung  $r$ ,  $p_\pi^{\text{old}}$
  - 2:  $\Lambda'_z = \{\lambda_i\}_{1:n} \leftarrow \text{top}_-(1/\eta)(\mathcal{H}, r)$
  - 3:  $\mathcal{S}_{\hat{\lambda}} \leftarrow \sum_{\Lambda'_z} w_i \cdot \hat{\lambda}(\lambda_i) \triangleright w_i = (n + 1) - i$
  - 4:  $\mathcal{S}_\pi \leftarrow \sum_{\Lambda'_z} w_i \cdot \pi(\lambda_i)$
  - 5:  $p_{\hat{\lambda}} \leftarrow p_\pi^{\text{old}} \cdot \mathcal{S}_{\hat{\lambda}} / (\mathcal{S}_\pi + \mathcal{S}_{\hat{\lambda}})$
  - 6:  $p_\pi \leftarrow p_\pi^{\text{old}} \cdot \mathcal{S}_\pi / (\mathcal{S}_\pi + \mathcal{S}_{\hat{\lambda}})$
  - 7: **return**  $p_\pi, p_{\hat{\lambda}}$
- 

#### 5.2.1 Decaying proportion of random sampling

Given the premise that we should initially trust the expert prior, yet with the benefit of incorporating random sampling (from Section 4), we make two additional assumptions, namely (i) we trust the expert’s belief most at the maximum fidelity  $z_{\max}$ ; and (ii) we would like to use cheaper fidelities to explore more. Given HB’s discretization of the fidelity range  $[z_{\min}, z_{\max}]$  into rungs  $r \in \{0, \dots, r_{\max}\}$ , we geometrically increase our sampling probability from  $\pi(\cdot)$  over  $\mathcal{U}(\cdot)$  by

$$p_\pi = \eta^r \cdot p_U, \quad (2)$$

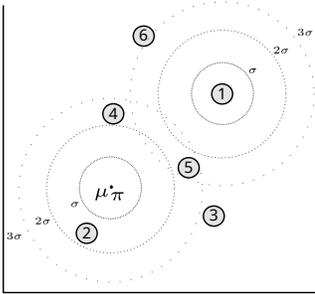


Figure 4: A visual example of how configurations contribute to the scores  $\mathcal{S}_\pi$  and  $\mathcal{S}_{\hat{\lambda}}$ . The prior distribution here is Gaussian  $\mathcal{N}(\mu_\pi, \sigma^2)$ , with a matching distribution  $\hat{\lambda}(\cdot) = \mathcal{N}(\hat{\lambda}, \sigma^2)$  placed on the incumbent. The labels 1-6 indicate configurations and their rank in  $\Lambda'_z$ . Here, e.g.,  $\lambda_2$  has densities such that  $\pi(\lambda_2) > \hat{\lambda}(\lambda_2)$ , contributing more to  $\mathcal{S}_\pi$  than to  $\mathcal{S}_{\hat{\lambda}}$ . The configurations  $\lambda_3, \lambda_4, \lambda_5$  are between both distributions and  $\lambda_6$  has low density under either, whereas  $\lambda_1$  (the incumbent) will be the primary influence such that  $\mathcal{S}_{\hat{\lambda}} > \mathcal{S}_\pi$ . This implies that  $p_{\hat{\lambda}} > p_\pi$  and demonstrates that all else being relatively equal, we will be more likely to sample from  $\hat{\lambda}(\cdot)$ .

with the constraint that  $p_\mu + p_\pi = 1$  (see L3, Algorithm 1). This naturally captures our assumptions, equally favouring  $\pi(\cdot)$  and  $\mathcal{U}(\cdot)$  initially but increasing trust in  $\pi(\cdot)$  according to the rung  $r$  we sample at. This geometric decay was inspired by similar scaling used in HB’s scheduling and backed by ablations over a linear decay and constant proportion (Appendix E.2.1).

## 5.2.2 Incumbent-based sampling proportion

Incumbent-based sampling intends to maintain strong anytime performance even under bad, uninformative, or adversarial priors. In `PriorBand`, initially  $p_{\hat{\lambda}} = 0$  until both (i) a budget equivalent to the first SH bracket has been exhausted ( $\approx \eta \cdot z_{\max}$ ); and (ii) at least one configuration has been evaluated at  $z_{\max}$ . This is captured in Line 4 of Algorithm 1 by `activate_incumbent()`.

Once incumbent-based sampling is active we need to decide how much to trust  $\hat{\lambda}(\cdot)$  vs. trusting  $\pi(\cdot)$ . Essentially, we achieve this by rating how likely the best seen configurations would be under  $\hat{\lambda}(\cdot)$  and  $\pi(\cdot)$ , respectively. Algorithm 2 shows the detailed steps to calculate  $p_\pi$  and  $p_{\hat{\lambda}}$ , and Figure 4 provides an example. We first define an ordered set  $\Lambda'_z = \{\lambda_1, \dots, \lambda_n\}$ , ordered by performance, of the top-performing  $(1/\eta)$  configurations<sup>1</sup> for the highest rung  $r$  with at least  $\eta$  evaluated configurations (Line 2 in Algorithm 2). Given  $\Lambda'_z$ , we compute two weighted scores  $\mathcal{S}_\pi, \mathcal{S}_{\hat{\lambda}}$ , capturing how likely these top configurations are under the densities of  $\pi(\cdot)$  and  $\hat{\lambda}(\cdot)$ , respectively. We also weigh top-performing configurations more highly, which is accomplished with a simple linear mapping  $w_i = (n + 1) - i$  (Lines 3-4 in Algorithm 2). Finally, we obtain  $p_\pi$  and  $p_{\hat{\lambda}}$  by normalizing these scores as a ratio of their sum (Lines 5-6 in Algorithm 2). We note the rates are adaptive based on how the set  $\Lambda'_z$  is spread out relative to the prior and the incumbent. We observe that this adaptive behavior is crucial in quickly recovering from bad prior inputs and thus maintaining strong final performance.

## 6 Experimental setup

We now describe our experiment design, benchmarks, and baselines and demonstrate the robustness of `PriorBand` in handling varying qualities of an expert prior input. Additionally, we also showcase the generality and practicality of the ESP,  $\mathcal{E}_\pi$ .

### 6.1 Benchmarks

We curated a set of 12 benchmarks that cover a diverse set of search spaces, including mixed-type spaces and log-scaled hyperparameters, and a wide range of downstream tasks, e.g., language modeling, image classification, tabular data, a medical application, and translation. We select 4 of the PD1 benchmarks (4 HPs) [32] that train large models such as transformers with batch sizes commonly found on modern hardware, and fit surrogates on them. Further, we select 5 benchmarks from LCBench (7 HPs) [33, 34] and consider all 3 JAHSBench [35] surrogate benchmarks that offer a 14 dimensional mixed-type search space for tuning both the architecture and training hyperparameters. All benchmarks and their selection are described in further detail in Appendix D.1.

<sup>1</sup>We want to use at least  $\eta$  configurations; hence  $\max(\eta, \# \text{ of configurations in the rung} / \eta)$  are selected.

## 6.2 Baselines

We choose a representative from each of the families of optimizers listed in Table 1, leaving out grid search. We use the official implementation of BOHB, while all other algorithms were implemented by us (and are available as part of our repository). The prior-based baselines ( $\pi$ BO and RS+Prior) sample the mode of the prior distribution as the first evaluation in our experiments to ensure a fair comparison where the prior configuration is certainly evaluated, irrespective of it being good or bad. For all the HB based algorithms we use  $\eta = 3$  and the fidelity bounds ( $z_{\min}, z_{\max}$ ) as per the benchmark. Further implementation and hyperparameter details for the baselines can be found in Appendix D.2. In principle, `PriorBand` only needs an additional input of  $\pi(\cdot)$  as the user belief, other than the standard HB hyperparameters. However, for a discussion on the hyperparameters of the incumbent-based local search (Section 5.1), please refer to Appendix E.2.5.

## 6.3 Design and setup

We show experiments both for single and multi-worker cases (4 workers in our experiments). For the single workers, we report the mean validation error with standard error bars for 50 seeds; for multi-worker experiments, we use 10 seeds. The plots show average relative ranks achieved aggregated across all 12 benchmarks when comparing the anytime incumbent configuration at  $z_{\max}$ . We group the runs on benchmark under different qualities of an expert prior input. We also compare the average normalized regret per benchmark under good priors. The prior construction procedure follows the design from the work by Hvarfner et al. [28], and we describe our procedure in detail in Appendix D.3. In the main paper, we only evaluate the good and bad prior settings as we believe that this reflects a practical setting of prior qualities; in Appendix F we also evaluate the near-optimal prior settings and also evaluate all 3 prior settings over high and low-performance correlation problems. In Appendix F.5, we also report additional experiments for different budget horizons. Further experiment design details can be found in Appendices D.4 and D.5. As an example of potential post-hoc analysis possible with `PriorBand`, we show how the dynamic probabilities calculated for each of the sampling strategies in `PriorBand` can be visualized (Appendix E.3), revealing the quality of the prior used.

# 7 Results

We now report and discuss the results of our experiments.

## 7.1 Robustness of `PriorBand`

Here, we demonstrate the robustness of `PriorBand` over a wide range of prior qualities by comparing them to the nearest non-prior-based algorithms: random search (RS) and `HyperBand` (HB). Figure 5(top) showcases `PriorBand` to be anytime equal or better than HB on each of our benchmarks under our good prior design. Note that the quality of a good prior varies per benchmark. The aggregation of Figure 5(top) is Figure 5(bottom-middle), which illustrates that `PriorBand` ( $\text{HB} + \mathcal{E}_\pi$ ) can utilize good priors and gain strong anytime performance over HB. Moreover, Figure 5(bottom-right) clearly shows `PriorBand`'s ability to recover from bad prior information and match vanilla-HB's performance in this adversarial case (the bad prior was intentionally chosen to be an extremely poor configuration – the worst of  $50k$  random samples). In most practical scenarios, the expert has better intuition than this setup and one can expect substantial speedups over HB, like in Figure 5(bottom-middle). Figure 5(bottom-left) demonstrates that when using an unknown quality of prior, even including the adversarial prior, `PriorBand` is still the best choice on average. We show a budget of 12 function evaluations here, which is approximately the maximal budget required for completing at least one HB iteration for the benchmarks chosen. In Appendix F.1 we show more comparisons to prior-based baselines and highlight `PriorBand`'s robustness.

## 7.2 Generality of Ensemble Sampling Policy $\mathcal{E}_\pi$

The ESP,  $\mathcal{E}_\pi$ , only needs access to an SH-like optimization state as input. In this section, we show that other popular multi-fidelity algorithms, such as ASHA or asynchronous-HB [36] that build on SH, can also support expert priors with the help of  $\mathcal{E}_\pi$ . In Figure 6, we compare the vanilla algorithms with their  $\mathcal{E}_\pi$ -augmented versions. Given that these algorithms were designed for parallel setups,

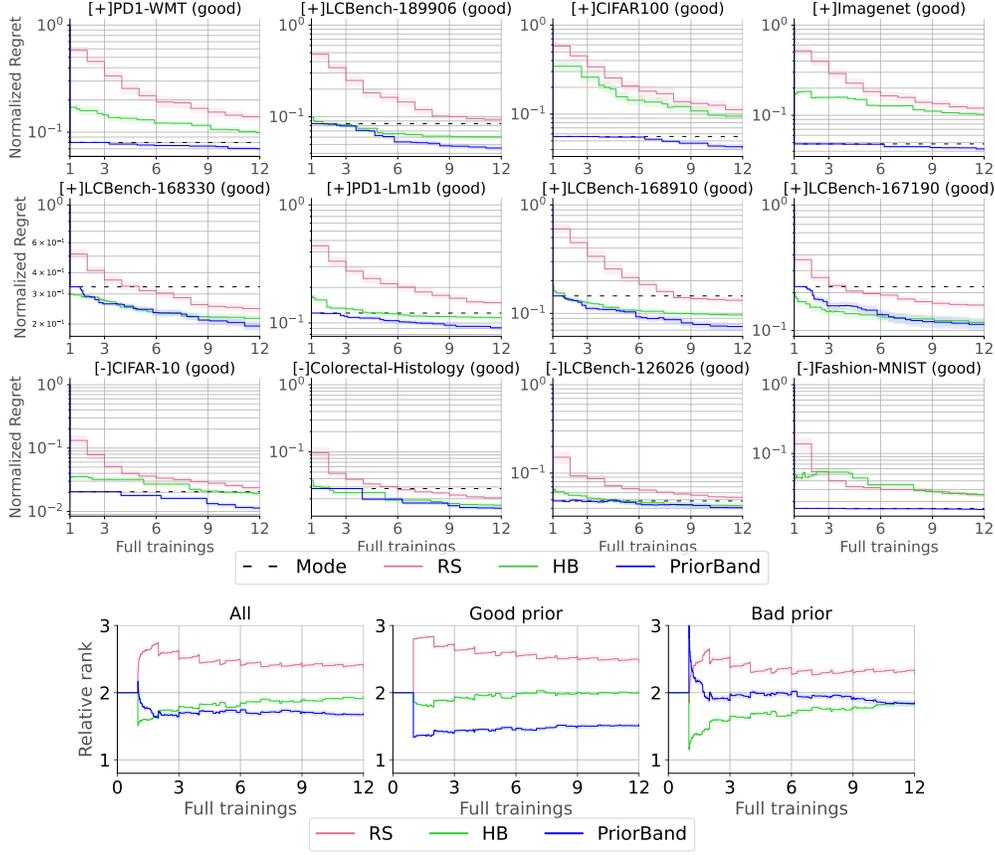


Figure 5: [Top] Comparing normalized regret in the good prior setting; The [Bottom] figure *Good* (middle) is a ranking aggregate view of the [Top] figure. [Bottom] Comparing average relative ranks of PriorBand to Random Search (RS) and HB, over single worker runs across benchmarks. Each benchmark-algorithm pair was run for 50 seeds where priors per benchmark are the same across a seed. We show mean rank and standard error. The *All* (left) plot averages the benchmark across Good and Bad priors.

we compare them on runs distributed among 4 workers, running for a total budget of 20 function evaluations. Similar to the previous section, the ESP-augmented (+ESP) algorithms can leverage good priors and recover under bad priors. Under bad priors, asynchronous-HB (+ESP) starts worst owing to bad prior higher fidelity evaluations at the start but shows strong recovery. In Appendix F.2 we compare these variants over different correlation settings of the benchmarks.

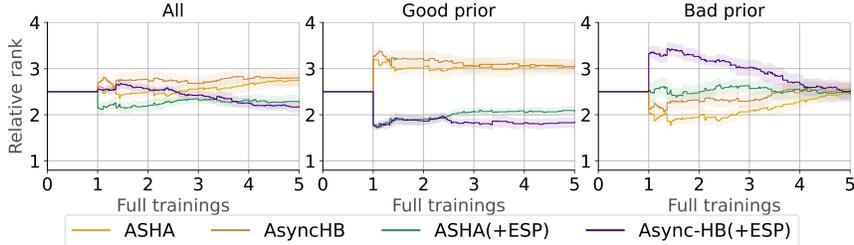


Figure 6: Comparing the average relative ranks of Asynchronous-SH (ASHA) and Asynchronous-HB (Async-HB) and their variants with the Ensemble Sampling Policy (+ESP) when distributed over 4 workers for a cumulative budget of 20 function evaluations.

### 7.3 Extensibility with models

Although our focus is the model-free low compute regime, we also show how `PriorBand` can be optionally extended with model-based surrogates to perform Bayesian Optimization, especially when longer training budgets are available. We compare BO,  $\pi$ BO, BOHB, and `PriorBand` with its model-extension, `PriorBand+BO` (Gaussian Processes as surrogates and Expected Improvement [37] for acquisition). All BO methods use an initial design of 10 function evaluations, except BOHB which sets this implicitly as  $N_{\text{dim}} + 2$ . Compared to other prior-based algorithms (RS+Prior,  $\pi$ BO), `PriorBand+BO` is consistently robust on average (*All*) under all priors. The anytime performance gain of `PriorBand+BO` under good priors is evidence of ESP  $\mathcal{E}_\pi$ 's utility. We note that `PriorBand+BO` recovers quickly in the bad prior setting and even outperforms all other algorithms after as little a budget as 6 full trainings. The fact that, until  $10\times$  function evaluations, `PriorBand+BO` is actually just `PriorBand` and model search has not begun, highlights the effectiveness of ESP in adaptively trading off different sampling strategies and the initial strength of incumbent sampling. Appendices E.4, F.3 contain details on our modeling choices and more experimental results.

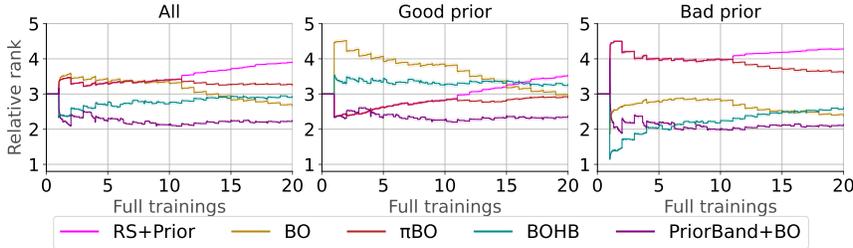


Figure 7: Comparing the average relative ranks of model-based algorithms and sampling from the prior (RS+Prior) under single-worker runs, aggregated over all benchmarks-per-prior quality.

### 7.4 Ablation studies

In Figure 8, we compare our choice of using density scores (Section 5.2.2) to trade-off prior-based and incumbent-based sampling against other similar or simpler heuristics, validating our decision. `PriorBand(constant)` employs a fixed heuristic to trade off prior and incumbent-based sampling as  $p_\pi = \eta \cdot p_\lambda$ . In `PriorBand(decay)`,  $p_\pi$  is decayed as a function of iterations completed. More precisely,  $p_\lambda = 2^b \cdot p_\pi$ , subject to  $p_\pi + p_\lambda + p_U = 1$  (from Section 5.2), where  $b \in \mathbb{N}$  and indicates the index of the current SH being run in HB. Appendix E.2.2 contains more ablations.

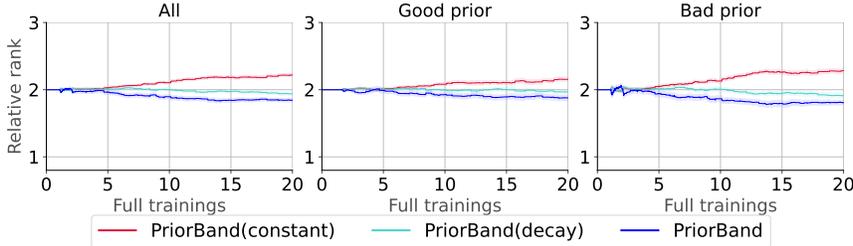


Figure 8: A comparison of `PriorBand` with 3 different strategies for trading off incumbent vs. prior distribution sampling. `PriorBand`'s default of using density scores shows a dominant performance in almost all scenarios. However, a bad prior induces a marked difference in performance. In general, the plot highlights that each variant reacts differently to bad priors. The adaptive nature of `PriorBand` clearly is robust to different scenarios comparatively.

## 8 Related work

While using expert priors and local search [38] for hyperparameter optimization has been explored previously, few works considered priors over the optimum [27, 28, 39], and they all target the

single-fidelity setting. The expert priors we consider should not be confused with the priors natively supported by BO, i.e., priors over the function structure determined by kernels [7, 40, 41].

In deep learning, training epochs and dataset subsets [17, 42, 43] are frequently used as fidelity variables to create cheap proxy tasks, with input resolution, network width, and network depth also occasionally used [35]. SH [29] and HB [30] are effective randomized policies for multi-fidelity HPO that use early stopping of configurations on a geometric spacing of the fidelity space and can also be extended to the model-based setting [44].

## 9 Limitations

We acknowledge that `PriorBand` inherits pathologies of HB such as poor fidelity correlations and can be sensitive to the choice of fidelity bounds and early stopping rate. Despite that, our results across different correlation settings suggest a relatively strong performance of `PriorBand` (Appendix F). Though `PriorBand` supports any kind of prior distribution as input, in our experiments we only considered the Gaussian distribution (with a fixed standard deviation of 0.25), as it is a natural choice and was used previously in the literature [28]. However, the expert is free to represent the prior with other distributions. We expect `PriorBand` to show similar behaviors as we report when comparing to other prior-based algorithms under a similar distribution. We note that `PriorBand` is not entirely free of hyperparameters. In our experiments, we keep all `PriorBand` hyperparameters fixed to remove confounding factors while comparing different prior strengths and correlations. Moreover, our hyperparameter choices for the experiments, such as  $\eta = 3$  and Gaussian priors are largely borrowed from existing literature [28, 36, 44]. Depending on the context of the specific deep learning expert, the compute budgets ( $\sim 10 - 12$  full trainings) used in our experimental setting may not always be feasible. However, the key insight from the experiments is that `PriorBand` can interface informative priors, and provide strong anytime performance under short HPO budgets (also in less than 5 full trainings). Longer HPO budgets for `PriorBand` ensure recovery from potential bad, uninformative priors. In practice, deep learning experts often have good prior knowledge and thus `PriorBand` is expected to retain strong anytime performance under low compute.

## 10 Conclusion

We identify that existing HPO algorithms are misaligned with deep learning (DL) practice and make this explicit with six desiderata (Section 1). To overcome this misalignment, our solution, `PriorBand`, allows a DL expert to incorporate their intuition of well-performing configurations into multi-fidelity HPO and thereby satisfies all desiderata to be practically useful for DL. The key component in `PriorBand`, the ESP,  $\mathcal{E}_\pi$ , is modular, flexible, and can be applied to other multi-fidelity algorithms.

## 11 Acknowledgments and disclosure of funding

N. Mallik, D. Stoll and F. Hutter acknowledge funding by the European Union (via ERC Consolidator Grant DeepLearning 2.0, grant no. 101045765). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. E. Bergman was partially supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215. C. Hvarfner and L. Nardi were partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. L. Nardi was supported in part by affiliate members and other supporters of the Stanford DAWN project — Ant Financial, Facebook, Google, Intel, Microsoft, NEC, SAP, Teradata, and VMware. Luigi Nardi was partially supported by the Wallenberg Launch Pad (WALP) grant Dnr 2021.0348. M. Lindauer was funded by the European Union (ERC, “ixAutoML”, grant no.101041029). The research of Maciej Janowski was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant number 417962828 and grant INST 39/963-1 FUGG (bwForCluster NEMO). In addition, Maciej Janowski acknowledges the support of the BrainLinks- BrainTools Center of Excellence.



## References

- [1] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [2] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *Proceedings of the International Conference on Learning Representations (ICLR'19)*, 2019. Published online: [iclr.cc](https://arxiv.org/abs/1812.06855).
- [3] Y. Chen, A. Huang, Z. Wang, I. Antonoglou, J. Schrittwieser, D. Silver, and N. de Freitas. Bayesian optimization in alphago. *arXiv:1812.06855 [cs.LG]*, 2018.
- [4] G. Melis, C. Dyer, and P. Blunsom. On the state of the art of evaluation in neural language models. In *Proceedings of the International Conference on Learning Representations (ICLR'18)*, 2018. Published online: [iclr.cc](https://arxiv.org/abs/1804.01902).
- [5] B. Zhang, R. Rajan, L. Pineda, N. Lambert, A. Biedenkapp, K. Chua, F. Hutter, and R. Calandra. On the importance of Hyperparameter Optimization for model-based reinforcement learning. In A. Banerjee and K. Fukumizu, editors, *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS'21)*, pages 4015–4023. Proceedings of Machine Learning Research, 2021.
- [6] R. Wightman, H. Touvron, and H. Jégou. Resnet strikes back: An improved training procedure in timm. In Ranzato et al. [45].
- [7] J. Snoek, H. Larochelle, and R. Adams. Practical Bayesian optimization of machine learning algorithms. In P. Bartlett, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, editors, *Proceedings of the 25th International Conference on Advances in Neural Information Processing Systems (NeurIPS'12)*, pages 2960–2968. Curran Associates, 2012.
- [8] M. Feurer, K. Eggenberger, S. Falkner, M. Lindauer, and F. Hutter. Auto-Sklearn 2.0: Hands-free automl via meta-learning. *JMLR*, 23(261):1–61, 2022.
- [9] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. In Larochelle et al. [46], pages 1877–1901.
- [10] Lambda. Demystifying gpt-3. <https://lambdalabs.com/blog/demystifying-gpt-3>, June 2020. Accessed on May 4, 2023.
- [11] X. Bouthillier and G. Varoquaux. Survey of machine-learning experimental methods at NeurIPS2019 and ICLR2020. Research report [hal-02447823], Inria Saclay Ile de France, 2020.
- [12] N. Hasebrook, F. Morsbach, N. Kannengießer, M. Zöller, J. Franke, M. Lindauer, F. Hutter, and A. Sunyaev. Practitioner motives to select hyperparameter optimization methods. *arXiv preprint arXiv:2203.01717*, 2022.
- [13] F. Schneider, Z. Nado, N. Agarwal, G. E. Dahl, and P. Hennig. Hity workshop poll, neurips 2022. <https://github.com/fsschneider/HITYWorkshopPoll>, 2022.
- [14] K. Van der Blom, A. Serban, H. Hoos, and J. Visser. Automl adoption in ml software. In K. Eggenberger, F. Hutter, E. LeDell, M. Lindauer, G. Shala, J. Vanschoren, and C. White, editors, *ICML workshop on Automated Machine Learning (AutoML workshop 2021)*, 2021.
- [15] Javier Gonzalez, Zhenwen Dai, Philipp Hennig, and Neil Lawrence. Batch bayesian optimization via local penalization. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 648–657, Cadiz, Spain, 09–11 May 2016. PMLR. URL <https://proceedings.mlr.press/v51/gonzalez16a.html>.

- [16] A. Klein, S. Falkner, S. Bartels, P. Hennig, and F. Hutter. Fast Bayesian optimization of machine learning hyperparameters on large datasets. In A. Singh and J. Zhu, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics (AISTATS'17)*, volume 54. Proceedings of Machine Learning Research, 2017.
- [17] K. Swersky, J. Snoek, and R. Adams. Freeze-thaw Bayesian optimization. *arXiv:1406.3896 [stats.ML]*, 2014.
- [18] T. Domhan, J. Springenberg, and F. Hutter. Speeding up automatic Hyperparameter Optimization of deep neural networks by extrapolation of learning curves. In Q. Yang and M. Wooldridge, editors, *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, pages 3460–3468, 2015.
- [19] C. White, A. Zela, B. Ru, Y. Liu, and F. Hutter. How powerful are performance predictors in neural architecture search? In Ranzato et al. [45].
- [20] M. Abdelfattah, A. Mehrotra, L. Dudziak, and N. Lane. Zero-cost proxies for lightweight NAS. In *Proceedings of the International Conference on Learning Representations (ICLR'21)*, 2021.
- [21] V. Nguyen and M. A. Osborne. Knowing the what but not the where in bayesian optimization. In H. Daume III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning (ICML'20)*, volume 98, pages 7317–7326. Proceedings of Machine Learning Research, 2020.
- [22] M. Feurer, J. Springenberg, and F. Hutter. Initializing Bayesian Hyperparameter Optimization via meta-learning. In B. Bonet and S. Koenig, editors, *Proceedings of the Twenty-ninth AAAI Conference on Artificial Intelligence (AAAI'15)*, pages 1128–1135. AAAI Press, 2015.
- [23] V. Perrone and H. Shen. Learning search spaces for bayesian optimization: Another view of hyperparameter transfer learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alche Buc, E. Fox, and R. Garnett, editors, *Proceedings of the 32nd International Conference on Advances in Neural Information Processing Systems (NeurIPS'19)*, pages 12751–12761. Curran Associates, 2019.
- [24] M. Feurer, B. Letham, F. Hutter, and E. Bakshy. Practical transfer learning for bayesian optimization. *arXiv:1802.02219v3 [stat.ML]*, 2018.
- [25] M. Wistuba and J. Grabocka. Few-shot bayesian optimization with deep kernel surrogates. In *Proceedings of the International Conference on Learning Representations (ICLR'21)*, 2021.
- [26] D. Stoll, J. Franke, D. Wagner, S. Selg, and F. Hutter. Hyperparameter transfer across developer adjustments. In Larochelle et al. [46].
- [27] A. Souza, L. Nardi, L. Oliveira, K. Olukotun, M. Lindauer, and F. Hutter. Bayesian optimization with a prior for the optimum. In N. Oliver, F. Pérez-Cruz, S. Kramer, J. Read, and J. Lozano, editors, *Machine Learning and Knowledge Discovery in Databases (ECML/PKDD'21)*, volume 12975 of *Lecture Notes in Computer Science*, pages 265–296. Springer-Verlag, 2021.
- [28] C. Hvarfner, D. Stoll, A. Souza, L. Nardi, M. Lindauer, and F. Hutter.  $\pi$ BO: Augmenting Acquisition Functions with User Beliefs for Bayesian Optimization. In *Proceedings of the International Conference on Learning Representations (ICLR'22)*, 2022. Published online: [iclr.cc](https://arxiv.org/abs/2202.08401).
- [29] K. Jamieson and A. Talwalkar. Non-stochastic best arm identification and Hyperparameter Optimization. In A. Gretton and C. Robert, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics (AISTATS'16)*, volume 51. Proceedings of Machine Learning Research, 2016.
- [30] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar. Hyperband: A novel bandit-based approach to Hyperparameter Optimization. *JMLR*, 18(185):1–52, 2018.
- [31] K. Kandasamy, G. Dasarathy, J. Schneider, and B. Póczos. Multi-fidelity Bayesian Optimisation with Continuous Approximations. In D. Precup and Y. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning (ICML'17)*, volume 70, pages 1799–1808. Proceedings of Machine Learning Research, 2017.

- [32] Zi Wang, George E. Dahl, Kevin Swersky, Chansoo Lee, Zelda Mariet, Zachary Nado, Justin Gilmer, Jasper Snoek, and Zoubin Ghahramani. Pre-trained gaussian processes for bayesian optimization, 2021. URL <https://arxiv.org/abs/2109.08215>.
- [33] Lucas Zimmer, Marius Lindauer, and Frank Hutter. Auto-pytorch tabular: Multi-fidelity metalearning for efficient and robust autodl. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):3079 – 3090, 2021.
- [34] F. Pfisterer, L. Schneider, J. Moosbauer, M. Binder, and B. Bischl. YAHPO Gym – design criteria and a new multifidelity benchmark for hyperparameter optimization. *arXiv:2109.03670 [cs.LG]*, 2021.
- [35] A. Bansal, D. Stoll, M. Janowski, A. Zela, and F. Hutter. JAHS-bench-201: A foundation for research on joint architecture and hyperparameter search. In *Proceedings of the 36th International Conference on Advances in Neural Information Processing Systems (NeurIPS’22)*, 2022.
- [36] L. Li, K. G. Jamieson, A. Rostamizadeh, E. Gonina, J. Ben-tzur, M. Hardt, B. Recht, and A. Talwalkar. A system for massively parallel hyperparameter tuning. In I. Dhillon, D. Papailiopoulos, and V. Sze, editors, *Proceedings of Machine Learning and Systems 2*, volume 2, 2020.
- [37] D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black box functions. *JGO*, 13:455–492, 1998.
- [38] Q. Wu, C. Wang, and S. Huang. Frugal optimization for cost-related hyperparameters. In Q. Yang, K. Leyton-Brown, and Mausam, editors, *Proceedings of the Thirty-Fifth Conference on Artificial Intelligence (AAAI’21)*. Association for the Advancement of Artificial Intelligence, AAAI Press, 2021.
- [39] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl. Algorithms for hyper-parameter optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, editors, *Proceedings of the 24th International Conference on Advances in Neural Information Processing Systems (NeurIPS’11)*, pages 2546–2554. Curran Associates, 2011.
- [40] K. Swersky, J. Snoek, and R. Adams. Multi-task Bayesian optimization. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, editors, *Proceedings of the 26th International Conference on Advances in Neural Information Processing Systems (NeurIPS’13)*, pages 2004–2012. Curran Associates, 2013.
- [41] C. Oh, E. Gavves, and M. Welling. BOCK : Bayesian optimization with cylindrical kernels. In *International Conference on Machine Learning*, pages 3865–3874, 2018.
- [42] A. Klein, S. Falkner, S. Bartels, P. Hennig, and F. Hutter. Fast bayesian Hyperparameter Optimization on large datasets. *Electronic Journal of Statistics*, 11:4945–4968, 2017.
- [43] V. Nguyen, S. Schulze, and M. A. Osborne. Bayesian optimization for iterative learning. In Larochelle et al. [46].
- [44] S. Falkner, A. Klein, and F. Hutter. BOHB: Robust and efficient Hyperparameter Optimization at scale. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning (ICML’18)*, volume 80, pages 1437–1446. Proceedings of Machine Learning Research, 2018.
- [45] M. Ranzato, A. Beygelzimer, K. Nguyen, P. Liang, J. Vaughan, and Y. Dauphin, editors. *Proceedings of the 34th International Conference on Advances in Neural Information Processing Systems (NeurIPS’21)*, 2021. Curran Associates.
- [46] H. Larochelle, M. Ranzato, R. Hadsell, M.-F. Balcan, and H. Lin, editors. *Proceedings of the 33rd International Conference on Advances in Neural Information Processing Systems (NeurIPS’20)*, 2020. Curran Associates.
- [47] L. Zimmer, M. Lindauer, and F. Hutter. Auto-Pytorch: Multi-fidelity metalearning for efficient and robust AutoDL. *TPAMI*, pages 3079–3090, 2021.

- [48] A. Falkner, G. Friedrich, K. Schekotihin, R. Taupe, and E. Teppan. Industrial applications of answer set programming. *KI-Künstliche Intelligenz*, pages 1–12, 2018.
- [49] X. Dong, M. Tan, A. Yu, D. Peng, B. Gabrys, and Q. Le. Autohas: Differentiable hyperparameter and architecture search. *arXiv:2006.03656 [cs.CV]*, 2020.
- [50] Justin M. Gilmer, George E. Dahl, and Zachary Nado. init2winit: a jax codebase for initialization, optimization, and tuning research, 2021. URL <http://github.com/google/init2winit>.
- [51] Y. Nesterov. A method of solving a convex programming problem with convergence rate  $O(1/\text{sqr}(k))$ . *Soviet Mathematics Doklady*, 27:372–376, 1983.
- [52] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 9:53–68, 2021.
- [53] Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, and Phillipp Koehn. One billion word benchmark for measuring progress in statistical language modeling. *CoRR*, abs/1312.3005, 2013. URL <http://arxiv.org/abs/1312.3005>.
- [54] Benjamin Lefaudeux, Francisco Massa, Diana Liskovich, Wenhan Xiong, Vittorio Caggiano, Sean Naren, Min Xu, Jieru Hu, Marta Tintore, Susan Zhang, Patrick Labatut, and Daniel Haziza. xformers: A modular and hackable transformer modelling library. <https://github.com/facebookresearch/xformers>, 2022.
- [55] Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. Findings of the 2015 workshop on statistical machine translation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 1–46, Lisbon, Portugal, September 2015. Association for Computational Linguistics. URL <http://aclweb.org/anthology/W15-3001>.
- [56] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [57] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [58] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the International Conference on Computer Vision and Pattern Recognition (CVPR’16)*, pages 770–778. CVF and IEEE CS, IEEE, 2016.
- [59] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [60] T. Chen and C. Guestrin. Xgboost: A scalable tree boosting system. In B. Krishnapuram, M. Shah, A. Smola, C. Aggarwal, D. Shen, and R. Rastogi, editors, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD’16)*, pages 785–794. ACM Press, 2016.
- [61] N. Awad, N. Mallik, and F. Hutter. DEHB: Evolutionary hyperband for scalable, robust and efficient Hyperparameter Optimization. In Z. Zhou, editor, *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI’21)*, pages 2147–2153, 2021.
- [62] J. Siems, L. Zimmer, A. Zela, J. Lukasik, M. Keuper, and F. Hutter. NAS-bench-301 and the case for surrogate benchmarks for Neural Architecture Search. *arXiv:2008.09777v4 [cs.LG]*, 2020.
- [63] B. Bischl, G. Casalicchio, M. Feurer, F. Hutter, M. Lang, R. Mantovani, J. N. van Rijn, and J. Vanschoren. Openml benchmarking suites and the openml100. *arXiv:1708.03731v1 [stat.ML]*, 2019.

- [64] J. Kiili, E. Laaksonen, R. Turner, D. Eriksson, S. Park, M. Mccourt, Z. Xu, and I. Guyon. Black box optimization challenge, 2020.
- [65] A. Klein, L. Tiao, T. Lienart, C. Archambeau, and M. Seeger. Model-based Asynchronous Hyperparameter and Neural Architecture Search. *arXiv:2003.10865 [cs.LG]*, 2020.
- [66] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020. URL <http://arxiv.org/abs/1910.06403>.
- [67] B. Matern. *Spatial Variation*. Lecture Notes in Statistics. Springer New York, 2013. ISBN 9781461578925. URL <https://books.google.se/books?id=HrbSBwAAQBAJ>.
- [68] Frank Hutter, Lin Xu, Holger H. Hoos, and Kevin Leyton-Brown. Algorithm runtime prediction: Methods & evaluation, 2013.

## A Resources used

All experiments in the paper were performed on cheap-to-evaluate surrogate benchmarks. We used several Intel(R) Xeon(R) Gold 6242 CPUs @ 2.80GHz to perform our experiments. Running one seed of one algorithm on one benchmark requires on average  $\sim 0.2$  core hours. For single-worker experiments, we run 21 algorithms with 50 seeds over 16 benchmarks with 3 strengths of priors per benchmark, totaling 50,400 single worker runs which equates to 10,080 core hours. In the parallel set of experiments, we use 4 cores per run, limiting ourselves to only 12 algorithms, 10 seeds, 7 benchmarks<sup>2</sup> with all 3 prior strengths. This totals 10,080 workers run in total which equates to 2,016 total core hours.

We additionally trained surrogate models for 2 metrics on 3 datasets for 4 hours with 8 cores, totaling another 192 core hours. During the development of our final algorithm, including failed experiments, re-runs, and preliminary testing, we estimate roughly another  $\sim 2,000$  core hours, a fifth of our total final cost. We estimate our total usage to have totaled  $\sim 14,288$  core hours.

## B Societal and environmental impact

Here, we discuss the potential societal and environmental impacts our work can have.

**Environmental** The contributed algorithm `PriorBand` and its re-usable component  $\mathcal{E}_\pi$  is designed to help reduce compute requirements for finding performant DL pipelines, thus reducing carbon emissions spent for HPO in DL. However, with the surplus compute available to many larger organizations, enabling robust methods for HPO could encourage further utilization of otherwise unused compute.

**Societal** Our paper and the contributed algorithm `PriorBand` are designed to assist a wide range of DL practitioners in finding performant hyperparameters. The ability of `PriorBand` to tune DL models under affordable compute enables practitioners to find strong hyperparameters otherwise only tenable for larger organizations. The societal impact depends on which task and DL pipeline `PriorBand` is applied to.

## C Licenses

- Our implementations - **MIT License**
- JAHS-Bench-201 benchmark [35] - **MIT License**
- YAHPO-Gym benchmark [34] - **Apache License 2.0**
- Learning curve benchmark [47] - **Apache License 2.0**
- PD1 [32] - **Apache License 2.0**
- BOHB [48] from `HpBandSter` - **BSD 3-Clause License**

## D Experiment details

### D.1 Benchmarks

Following Equation 1, we frame all the benchmark tasks as a minimizing problem. The benchmarks we use are provided by our curated suite of multi-fidelity benchmarks (`mf-prior-bench`) that treats priors as first-class citizen. We include our own synthetic Hartmann functions (D.1.1) extended to the multi-fidelity setting. We wrap JAHS-Bench-201 [35] (D.1.2) and Yahoo-Gym [34] (D.1.4) and provide new surrogate benchmarks for large models for image and language tasks, trained from optimization data obtained from the PD1 benchmark from HyperBO [32] (D.1.3).

---

<sup>2</sup>Please see D.1.6 as for why we do not include all benchmarks for parallel runs.

### D.1.1 Multi-fidelity synthetic Hartmann (MFH)

The multi-fidelity Hartmann functions follow the design of Kandasamy et al. [31], where, for  $[0, 1]$ -scaled  $z$ , the fidelity is parameterized as

$$g(\mathbf{x}, z) = \sum_{i=1}^4 (\alpha_i - \alpha'_i(z; b)) \exp \left( - \sum_{j=1}^D A_{ij} (x_j - P_{ij})^2 \right) \quad (3)$$

where for Hartmann-3,

$$A = \begin{bmatrix} 3 & 10 & 30 \\ 0.1 & 10 & 35 \\ 3 & 10 & 30 \\ 0.1 & 10 & 35 \end{bmatrix}, \quad P = 10^{-4} \times \begin{bmatrix} 3689 & 1170 & 2673 \\ 4699 & 4387 & 7470 \\ 1091 & 8732 & 5547 \\ 381 & 5743 & 8828 \end{bmatrix},$$

and for Hartmann-6,

$$A = \begin{bmatrix} 10 & 3 & 17 & 3.5 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{bmatrix}, \quad P = 10^{-4} \times \begin{bmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{bmatrix}$$

where  $\alpha'_i(z; b) = b(1 - z_i)$ . In the original paper, the variable which accounts for the bias between fidelities,  $b$ , is set to 0.1. To account for the fact that we only consider a single fidelity variable, we set  $z_i = z, \forall i$ , and increase the bias terms significantly to create realistic task correlations.

For the **good correlation** version used in our evaluations, we set  $b = 2.5$  with half-normally distributed noise of  $\sigma = 2(1 - z)$ .

The **bad correlation** version uses  $b = 4$  with a noise of  $\sigma = 5(1 - z)$ .

Tables 2, 3 show the search space for this synthetic benchmark.

name	type	values	info
X_0	continuous	[0.0, 1.0]	
X_1	continuous	[0.0, 1.0]	
X_2	continuous	[0.0, 1.0]	
z	log integer	[3, 100]	fidelity

Table 2: Synthetic Multi-Fidelity Hartmann search space in 3 dimensions.

name	type	values	info
X_0	continuous	[0.0, 1.0]	
X_1	continuous	[0.0, 1.0]	
X_2	continuous	[0.0, 1.0]	
X_3	continuous	[0.0, 1.0]	
X_4	continuous	[0.0, 1.0]	
X_5	continuous	[0.0, 1.0]	
z	log integer	[3, 100]	fidelity

Table 3: Synthetic Multi-Fidelity Hartmann search space in 6 dimensions.

The global minimums of these functions are known:

- Hartmann-3d:  $f(x^*) = -3.86278$  at  $x^* = (0.114614, 0.555649, 0.852547)$
- Hartmann-6d:  $f(x^*) = -3.32237$  at  $x^* = (0.20169, 0.150011, 0.476874, 0.275332, 0.311652, 0.6573)$

### D.1.2 JAHS-Bench-201

JAHS-Bench-201 [35] is a benchmark consisting of surrogates trained on 140 million data points of Neural Networks trained on 3 datasets, namely CIFAR10, Colorectal-Histology, and Fashion-MNIST. They extend the search space beyond the original tabular search space of NAS-Bench-201 [49] consisting of purely discrete architectural choices, introducing both hyperparameters and multiple fidelities to create the first multi-multi-fidelity benchmark for deep learning hyperparameter optimization (Table 4).

Each of the three datasets shares equal search spaces while we fix the fidelity parameters, depth  $N$  and width  $W$ , to their maximum. We further limit `Resolution` to a fixed value of 1.0 out of the three

original choices  $\{0.25, 0.5, 1.0\}$ . The surrogates provided by JAHS-Bench-201 do not explicitly model the monotonic constraint that as epoch increase, so should the training cost. In practice, this was found to be insignificant but we state so for completeness. For these benchmarks, optimizers minimize  $1 - \text{valid\_acc}$ .

name	type	values	info
Activation	categorical	{ReLU,Hardswish,Mish}	
LearningRate	continuous	[0.001, 1.0]	log
N	constant	5	
Op1	categorical	{0,1,2,3,4}	
Op2	categorical	{0,1,2,3,4}	
Op3	categorical	{0,1,2,3,4}	
Op4	categorical	{0,1,2,3,4}	
Op5	categorical	{0,1,2,3,4}	
Op6	categorical	{0,1,2,3,4}	
Optimizer	constant	SGD	
Resolution	constant	1.0	
TrivialAugment	categorical	{True,False}	
W	constant	16	
WeightDecay	continuous	[1e-05, 0.01]	log
epoch	integer	[3, 200]	fidelity

Table 4: The JAHS-Bench-201 search space for all 3 datasets, CIFAR10, Colorectal-Histology and Fashion-MNIST.

### D.1.3 PD1 (HyperBO)

The PD1 benchmarks consist of surrogates trained on the learning curves of large architectures, spanning both natural language and computer vision tasks. The original tabular data is obtainable from the output generated by HyperBO [32] using the dataset and training setup of [50], enabling us to test our methods and baselines for low-budget settings, where multi-fidelity methods are most applicable. The hyperparameters considered for the optimization runs were for Nesterov Momentum [51] which constitutes our search space (Table 5-8). All other hyperparameters were fixed according to their training setup and provided data.

This tabular data consists of 4 collections of optimization records, a grid-like spread of configurations and also those chosen by their optimizer, in both an initial testing phase and a later full experiment phase. To maximize the data available to the surrogate, we utilize all of this data but take care to drop duplicated runs from their test runs. We select these 4 benchmarks out of the available 24, opting to have a variety of tasks, favoring larger models where possible, or tasks that use big batch sizes such as 2048.

For these benchmarks, each optimizer aims to minimize the `valid_error_rate`. The hyperparameters listed are based on the minimum and maximum values found within the original tabular data, rather than the reported ranges by the authors of HyperBO [32]. This was to prevent surrogates from being required to extrapolate outside of their training domain.

- **lm1b\_transformer\_2048** derives from the optimization trace of a transformer model [52] with batch size of 2048 on the 11mb statistical language modelling dataset [53].
- **translatewmt\_xformer\_64** derives from the optimization trace of an xformer [54] with batch size of 64 on the WMT15 German-English [55].
- **cifar100\_wideresnet\_2048** derives from the optimization trace of a wideresnet model [56] with batch size 2048 on the cifar100 dataset [57].
- **imagenet\_resnet\_512** derives from the optimization trace of a resnet model [58] with batch size of 512 on the imagenet dataset [59].

name	type	values	info
lr_decay_factor	continuous	[0.010543, 0.9885653]	
lr_initial	continuous	[1e-05, 9.986256]	log
lr_power	continuous	[0.100811, 1.999659]	
opt_momentum	continuous	[5.9e-05, 0.9989986]	log
epoch	integer	[1, 74]	fidelity

Table 5: The `lm1b_transformer_2048` search space.

name	type	values	info
lr_decay_factor	continuous	[0.0100221257, 0.988565263]	
lr_initial	continuous	[1.00276e-05, 9.8422475735]	log
lr_power	continuous	[0.1004250993, 1.9985927056]	
opt_momentum	continuous	[5.86114e-05, 0.9989999746]	log
epoch	integer	[1, 19]	fidelity

Table 6: The `translatewmt_xformer_64` search space.

**Training surrogates on the PD1 tabular data** The original data is a mix of several datasets, models, and their parameters for which we do some preprocessing. All data-preprocessing is available as part of `mf-prior-bench` and consists of:

1. Splitting the raw data by all available  $\{datasetname, model, batchsize\}$  subsets.
2. Identify which columns are hyperparameters by those being marked as such and consist of more than one unique value.
3. Drop all columns which are not hyperparameters or metrics.
4. Drop all NaN values for which no metrics are recorded.
5. Drop configurations that recorded *divergent* training costs.
6. Drop duplicated entries, preferring to keep those from their full experimental runs.

To decide if a configuration diverged was to find outliers that reported extreme outlier costs, with a cutoff applied heuristically to each individual dataset, the details of which can be found within `mf-prior-bench`. This was done to ease the learning process of the surrogate model and to remove emphasis on these outliers. The resulting surrogate is no longer aware of these divergences and offers smooth interpolation for the training cost for these configurations.

Once the datasets are prepared, we then train a single surrogate XGBoost model [60] per metric recorded. This training was performed using DEHB [61], optimizing for the mean R2 loss of 5-fold cross-validation for a total of 4 hours, 8 CPU cores, and the seed set to 1. All surrogate models were found to converge in their R2 loss. These surrogates are available as part of `mf-prior-bench` for further inspection. While certainly improvements can be made in this modeling phase, for the purpose of our experiments, they offer a good approximation of the entire optimization landscape.

name	type	values	info
lr_decay_factor	continuous	[0.010093, 0.989012]	
lr_initial	continuous	[1e-05, 9.779176]	log
lr_power	continuous	[0.100708, 1.999376]	
opt_momentum	continuous	[5.9e-05, 0.998993]	log

Table 7: The `cifar100_wideresenet_2048` search space.

name	type	values	info
lr_decay_factor	continuous	[0.010294, 0.989753]	
lr_initial	continuous	[1e-05, 9.774312]	log
lr_power	continuous	[0.100225, 1.999326]	
opt_momentum	continuous	[5.9e-05, 0.998993]	log

Table 8: The imagenet\_resnet\_512 search space.

#### D.1.4 Yahpo-Gym

The Yahpo-Gym [34] collection is a large collection of multi-fidelity surrogates across a wide range of tasks, including traditional machine learning models with dataset size as a fidelity, as well as Neural Network, benchmarks such as LCBench and NAS-Bench-301 [62].

Yahpo-Gym provides surrogates trained on a shared search space between all of these tasks. We ignore the rest of the available benchmarks from Yahpo-Gym as the others consist of non-DL tasks or contain conditional search spaces which are not suitable for most of our baselines. For these benchmarks, the optimization objective was to minimize  $1 - \text{val\_balanced\_accuracy}$ . Table 9 shows the search space for these 5 benchmarks.

name	type	values	info
batch_size	integer	[16, 512]	log
learning_rate	continuous	[0.0001, 0.1]	log
max_dropout	continuous	[0.0, 1.0]	
max_units	integer	[64, 1024]	log
momentum	continuous	[0.1, 0.99]	
num_layers	integer	[1, 5]	
weight_decay	continuous	[1e-05, 0.1]	
epoch	integer	[1, 52]	fidelity

Table 9: The lcbench search space.

For our experiments, we choose 5 LCBench tasks from 34 from OpenML [63] according to the Spearman rank correlation of configurations at the 10% epoch and the final epoch  $z_{\max}$ . The chosen tasks, 126026, 167190, 168330, 168910, and 189906 are equally spaced according to this correlation and include the task featuring the least (126026) and most (189906) correlation. The correlation for all LCBench tasks can be seen in Figure 9.

#### D.1.5 Classifying benchmarks into high-low correlation

Our definition of a *high* or a *low* correlation benchmark is that the benchmark must have 0.8 spearman correlation of rankings at 10% of the maximum fidelity to be classified as *high*, otherwise it is classified as *low*. We depict this in Figure 10, showing the spearman correlation between each fidelity available and the final full fidelity. While other cutoffs and classifications are possible, given our suite of benchmarks, we find this to be a reasonable separation given the data.

To further motivate this choice, Figure 11 shows the performance of random search and HyperBand across these 12 benchmarks. We choose HyperBand as under the hood, it runs different instantiations of SH. Under a single worker-run, HB becomes a sequential run of SH with increasing  $z_{\min}$ . Strong performance of SH requires a high correlation of performance across fidelities. Thus in  $5\times$  in Figure 11, wherever HB the performance gap between random search and hyperband is not pronounced in the short budget regime shown, it can be inferred that the lower fidelities do not provide reliable prediction when early stopping configurations. Based on the classification strategy derived from Figure 10, we denote high correlation benchmarks with a (+) and low correlation benchmarks with a (−) in Figure 11. The comparison of random search and HB in this figure reflects the reasonable correctness of our classification strategy.

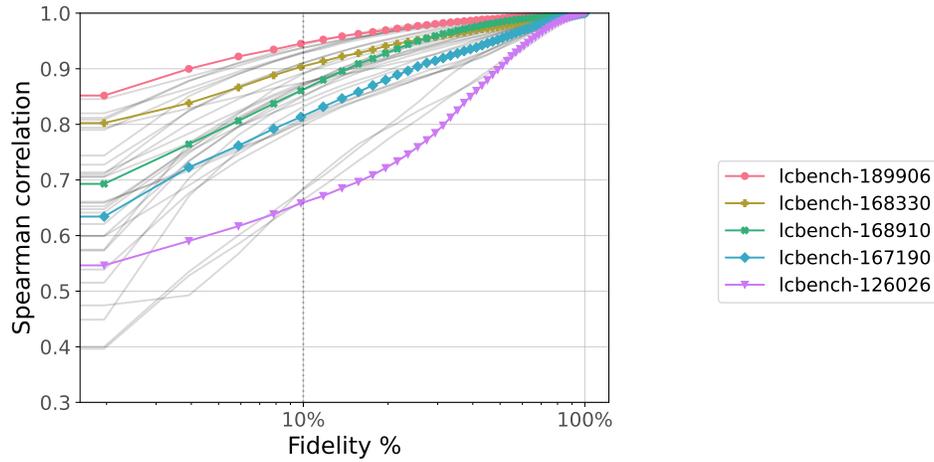


Figure 9: Each curve shows the spearman-correlation of 25 configurations from the given fidelity along the x-axis to the last fidelity, where the standard deviation is estimated with repeated samples until the mean curve converges to within a 0.001 Euclidean distance update to the previous mean. The highlighted lines are the LCBenchmark tasks selected by taking the (0, 0.25, 0.5, 0.75, 1) quantiles of the correlations at 10% of the full fidelity range. The faded gray lines with no markers represent LCBenchmark tasks not selected.

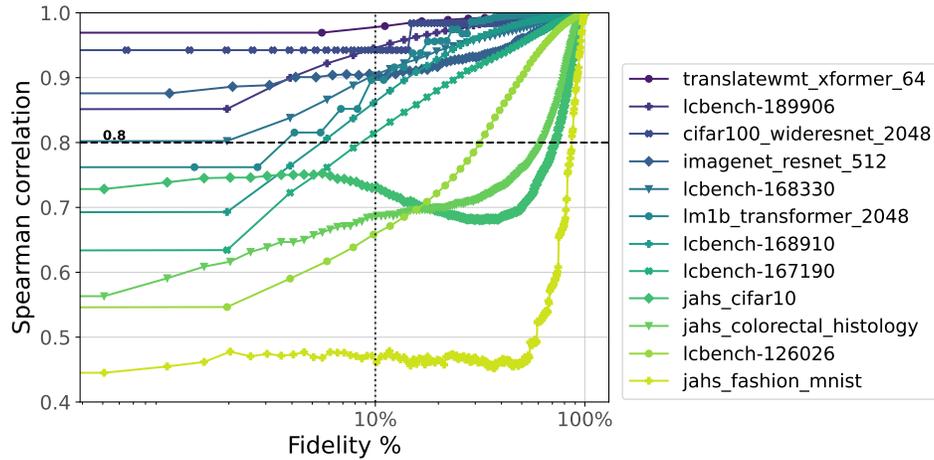


Figure 10: Each curve shows the spearman-correlation of 25 configurations from the given fidelity along the x-axis to the last fidelity, where the standard deviation is estimated with repeated samples until the mean curve converges to within a 0.001 Euclidean distance update to the previous mean. Our good/bad correlation definition corresponds to a spearman-correlation cutoff of 0.8 at 10% of the maximum budget. The legend is sorted by their correlation at 10% fidelity and corresponds to the order of benchmarks in Figure 11.

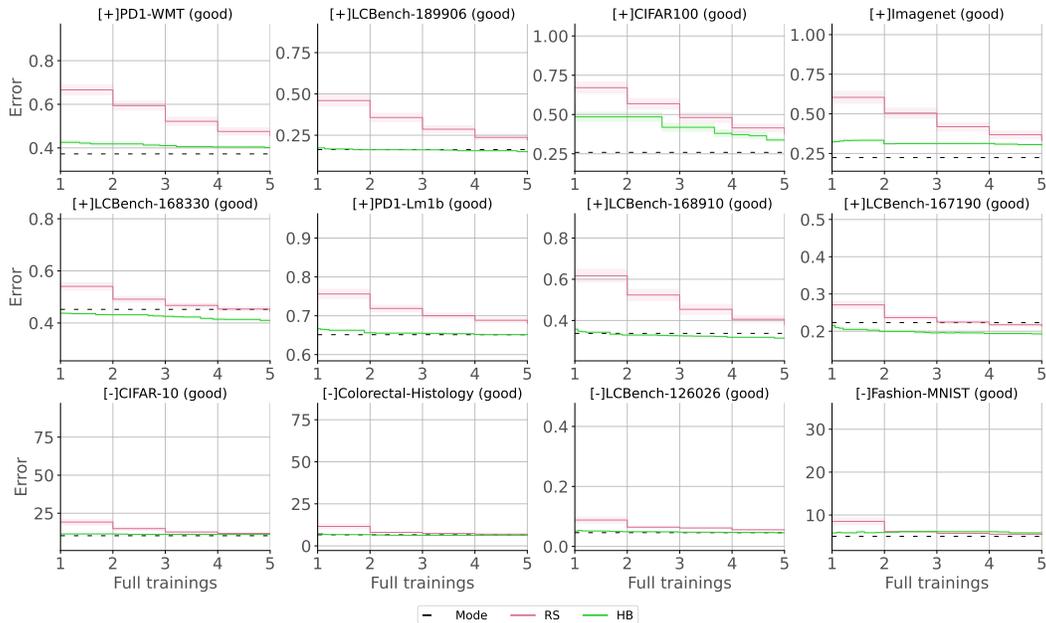


Figure 11: Comparing Random Search and HyperBand over  $5\times$  budget to gauge quality of performance correlation across fidelities. We classify benchmarks marked with  $[+]$  as *high correlation* benchmarks (Row 1 and 2). The benchmarks marked as  $[-]$  are classified as *low correlation* benchmarks (Row 3).

### D.1.6 Issues with Benchmarks in parallel setting

During our parallel worker runs, we noticed that workers running on LCBench (see D.1.4) would silently drop out. This is an artifact of Yaho-Gym, where the shared loading of system resources does not play nicely with workers being started in parallel. The optimization runs would still progress without issue but only utilize 1-3 processes instead of the deployed 4. The degree to which the problem occurred was minimal but likely to impact aggregated results. As a safety precaution, we remove LCBench from our parallel algorithm evaluations to prevent undue biases from leaking into our evaluations.

## D.2 Baselines

For fair comparison, customizability and certain technical constraints, we reimplement all the baselines listed below, other than BOHB.

**HyperBand** We implement our own version of HB[30] to allow for the input of priors. We verified our implementation with the popular HB implementation provided in BOHB [44]. We use  $\eta = 3$  for all experiments with the minimum and maximum budget coming as an input from the problem to solve, in this case, benchmarks. As described in Section 3, HB iterates over different instantiations of SuccessiveHalving (SH). In Figure 12 we show an example<sup>3</sup> for an SH run under  $\eta = 2$ .

**Bayesian optimization** is a prominent framework for Hyperparameter Optimization [7, 64], hence we choose its Gaussian Processes (GP) implementation as a model-based competitor. We incorporate expert priors to the optimization following the  $\pi$ BO algorithm [28]. In a low-budget setting, model-based search proves challenging for high-dimensional search spaces (e.g. JAHS-Bench-201) as common practices require the number of initial random observations equal to the search space dimensionality. To allow model-based search in BO and  $\pi$ BO we set their initial design size to 10, to allow model fitting under our experiment budgets. Expected Improvement is used as the acquisition function across all BO algorithms.

<sup>3</sup>image and caption sourced under CC-BY-4.0 from [https://www.automl.org/wp-content/uploads/2019/05/AutoML\\_Book.pdf](https://www.automl.org/wp-content/uploads/2019/05/AutoML_Book.pdf)

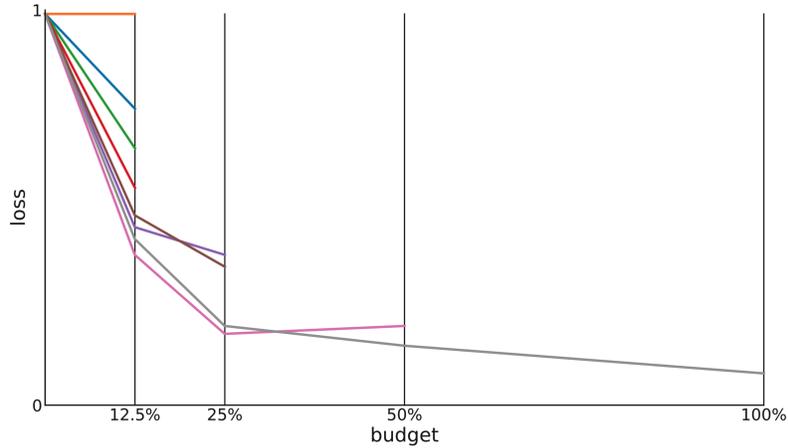


Figure 12: Illustration of `SuccessiveHalving` for eight algorithms/configurations. After evaluating all algorithms on 1/8 of the total budget, half of them are dropped and the budget given to the remaining algorithms are doubled.

**BOHB** [44]<sup>4</sup> incorporates multi-fidelity HB into the Bayesian optimization framework by building KDE models on each fidelity level to efficiently guide the search. The official implementation has no direct way of accepting a prior distribution over optimal configurations and incorporating it into search. We keep the other default settings intact.

**ASHA** [36] was developed as an extension to `SuccessiveHalving` that can run on massively parallel systems, designed to minimize idle workers. ASHA modified SH and HB to allow for a configuration to be promoted as soon as  $\eta$  configurations have been seen at a fidelity, calling such a step an *asynchronous promotion*. Thereby, a free worker need not remain idle till an entire predefined number of configurations have finished evaluation at a fidelity, like vanilla-SH or HB. Since HB effectively runs multiple SH brackets, asynchronous-HB can be designed as HB running multiple ASHA brackets. However, for this work, we chose the sampling distribution for the brackets as used in Klein et al. [65] for asynchronous-HB.

**Mobster** [65] extended asynchronous-HB with a surrogate model to create an asynchronous version of BOHB. The asynchronous design aside, Mobster is different from BOHB as it uses Gaussian Processes as the surrogate model, unlike BOHB which uses Tree Parzen Estimators (TPE).

**ESP based baselines.** We construct multiple baselines that do not explicitly exist in the literature but serve as important baselines for a fair comparison and a deeper understanding of the problem and the method.

- $[X]+Prior$  methods: for such a method, the uniform random search component in  $X$  is replaced with sampling from the prior. If accompanied by a  $Y\%$ , that indicates the percentage of prior-based sampling, with  $(1 - Y)\%$  for random sampling. For example,  $HB+Prior$  is HB but with only sampling from prior and  $HB+Prior(30\%)$  would indicate that sampling from prior happens with a probability of 0.3 and random sampling with a probability of 0.7.
- $[X]+ESP$  methods: for such a method, the uniform random search component in  $X$  is replaced with sampling from the ESP,  $\mathcal{E}_\pi$ .
- $PriorBand+BO$ : this runs `PriorBand` as described in Section 5 for a budget equivalent to an initial design of 10, and then switching to sampling a configuration from the acquisition in a BO loop, disabling the ESP.

**Parallel runs.** We run each benchmark-optimizer-seed combination over 4 workers. For HB based algorithms, if a worker is free, an evaluation from the next SH bracket is already started, in order to maximize worker efficiency. However, a pending evaluation from the earliest active SH bracket

<sup>4</sup>Implementation from `HpBandSter`: <https://github.com/automl/HpBandSter>

has the highest priority. For BO algorithms, under our parallel setup for multi-fidelity optimization, we did not require batch acquisitions. However, we needed to account for incomplete evaluations from configurations that are still training when fitting the surrogate. We do this simply by fitting the surrogate on the finished evaluations and predicting the mean for the pending evaluations, before retraining on the required set of configurations<sup>5</sup>.

### D.2.1 GP kernels

**Numerical kernel.** In the numerical continuous and discrete domain, we use the standard ([7, 66]) *Matérn-5/2* kernel [67]. For  $D$ -dimensional numerical inputs  $x$  and  $x'$ ,

$$k_{M5/2}(x, x') = \theta_0 \left( 1 + \sqrt{5r^2(x, x')} + \frac{5}{3}r^2(x, x') \right) \exp \left( -\sqrt{5r^2(x, x')} \right), \quad (4)$$

where  $r^2(x, x') = \sum_{d=1}^D (x_d - x'_d)^2 / \theta_d^2$  denotes a scaled Euclidean distance between points, and  $\theta_d$  is a dimension-specific lengthscale.

**Categorical kernel.** A straightforward extension of a Matérn kernel for a categorical domain is to use *1-in-K encoding*. However, this solution increases the dimensionality of the input, which might result in poor regression performance. Instead, we follow [68] who propose direct handling of categorical inputs by computing a weighted Hamming distance:

$$k_{HM}(x, x') = \exp \left( \sum_{d=1}^D \left( -\theta_d \cdot \mathbb{1}_{x_d \neq x'_d} \right) \right) \quad (5)$$

### D.3 Generating priors

We borrow the prior generation procedure in Hvarfner et al. [28] to generate the *near optimum* and *bad* priors. Additionally, we construct another class of *good* priors to reflect a different strategy that may be closer to practice, since for DL, the optimum is generally not known. We thus generate three kinds of priors, *near optimum*, *bad*, and *good*, the first two to simulate boundary conditions of priors that may be received and *good* to simulate a practitioner with some prior knowledge of a good configuration to choose. In each of these cases, the prior distribution is a normal  $\mathcal{N}(\lambda, \sigma^2)$ , with  $\sigma = 0.25$ , where the configuration  $\lambda$  is generated by the following processes:

**Near optimum:** Given we don't know the optimum configuration for a benchmark, we approximate this by taking the best of 50,000 configurations according to their observed loss value. During each seeded run, we perturb this configuration by Gaussian noise with a  $\sigma = 0.25$  for numerical values, while using this same 0.25 for uniformly selecting a different categorical value. This is done to ensure there is still some room for improvement possible over the initial prior configuration. All algorithms receive the same prior configuration as defined by the seed. In the case of the Multi Fidelity Hartmann benchmarks, we can access the optimum of the function analytically and so we take this optimum as the configuration which is to be perturbed.

**Good:** We define a *good* prior as one that is suggested by limited prior evaluations done by a practitioner. For this, we take the best of 25 configurations and apply no further noise modifications per seed. This means each run will see the same prior irrespective of the prior. We choose a budget of 25 random samples since in our experimental setup our focus is strictly on HPO budgets under  $20 \times$ . As we show later, the relative quality of this class of priors varies per benchmark, compared to the near-optimum priors.

**Bad:** As we don't know the worst configuration for each benchmark, we simulate this process by taking the worst configuration of 50,000 samples. No additional noise is added per seed, to ensure we always treat the worst known configuration as the prior input.

It is worth noting that while *near optimum* and *bad* are the best and worst of 50,000 samples, respectively, it need not be that these configurations are also the best and worst for search. For benchmarks with a high categorical count, such as JAHs Bench as described in D.1.2 (middle row

<sup>5</sup>unlike typical BO, multi-fidelity BO may have different data subsets to fit the surrogate over

in Figure 13), simply switching a categorical value can drastically alter the performance of the configuration. More generally, a single point of the search space does not immediately provide information about the performance in its locality. As such, bad performing configurations could be reached with the noise perturbations of the *near optimum* prior.

To investigate the performance impacts of each prior, in Figure 13 we plot the distribution of performances of 25 configurations sampled from the prior distribution used. This is additionally done for 50 seeds, giving a total of 1,250 configurations plotted per violin.

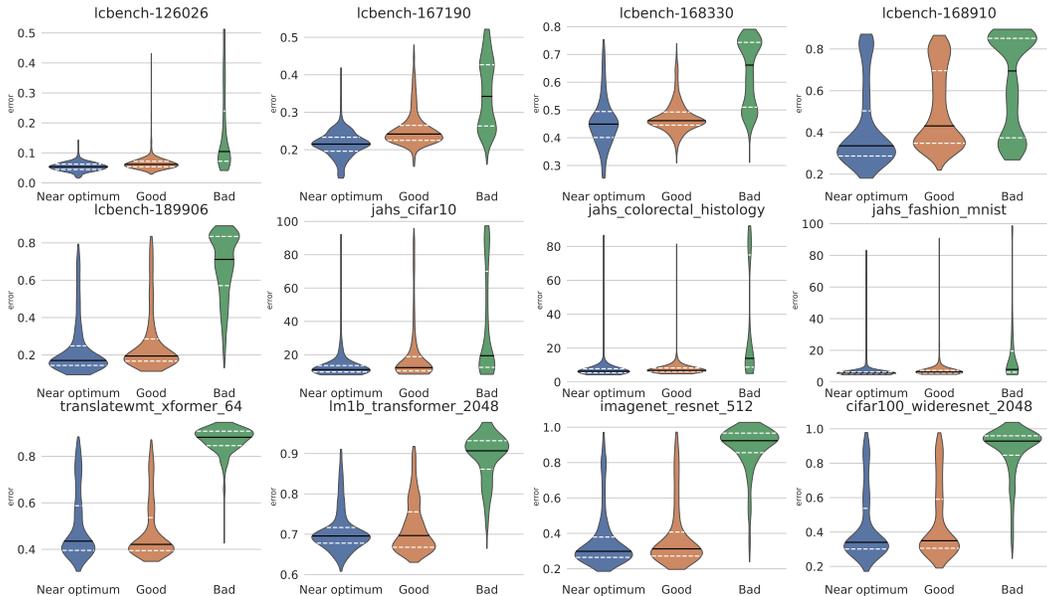


Figure 13: Each violin shows the density of performances achieved by random search on a distribution using the given prior. The black line in the violin is the mean with the dashed white lines indicating the 25% and 75% quartiles. In all cases, we see *near optimum* and *good* outperform the *bad* prior, as expected. In the case of the PD1 datasets (bottom row), some *good* priors achieve a better mean error than *near optimum* but their best configurations found are not as strong as they are when using *near optimum* priors. This could be a result of the near optimum prior being perturbed each run. Some iterations may put the prior at a slightly worse region than the fixed good prior. However, the near-optimum priors generally achieve the best-seen score.

#### D.4 Experimental setup

Generally, the HPO landscape for a new task is unknown and the quality of the expert prior input cannot be gauged without previous run data available. In order to simulate how an expert may use HPO in their task, we design and present the experiments to represent such scenarios. We show the aggregated results (*All*) over different qualities of prior input to highlight robustness across prior qualities, which may be unknown at the beginning of the problem. Subsequently, we break the results down into performances under the *near optimum* priors, *good* prior, and *bad* prior. This setup intends to highlight that `PriorBand`'s performance benefits are more substantial if the quality of the prior is better. Our experiment design substantiates the hypothesis that `PriorBand` and algorithms extended with  $\mathcal{E}_\pi$  are robust across any given prior input, a property not held by our baseline algorithms.

The metrics reported for single-worker runs are for 50 seeds across 12 benchmarks, 5 from LCBench (D.1.4), 3 from JAHS-Bench-201 (D.1.2) and 4 from PD1 (D.1.3). For the parallel case, we run 4 workers for each algorithm, 10 seeds, for only 7<sup>6</sup> of these benchmarks. We chose 10 seeds due to the increased cost of running many workers<sup>7</sup>.

<sup>6</sup>We exclude 5 LCBench benchmarks due to parallelism issues as described in D.1.6

<sup>7</sup>asynchronicity is simulated for the queries to the benchmarks by sleeping for *epoch* seconds

## D.5 Evaluation protocol

For the reduction factor  $\eta$ , in HB and thereby PriorBand, we chose  $\eta = 3$ . For the various  $[z_{\min}, z_{\max}]$  coming from the 12 benchmarks chosen, there are 2, 3, or 4 levels of fidelity created (rungs) in HB for our set of experiments. Therefore, for the single worker case, we show a budget of 12 full function evaluations, which is approximately equivalent to the budget of one complete HB iteration with 4 fidelity levels ( $12 \times z_{\max}$ ). For other cases, we show a budget of  $20\times$  and report under  $5\times$  for multi-worker runs.

The evaluation setup is motivated by how an HPO algorithm might be used in practice. The HPO algorithm stopped anytime, will return the best-seen configuration, i.e., will return the current incumbent. To compare algorithms, we plot the incumbent over the budget spent in epochs. The incumbent is simply chosen as the configuration with the lowest error across all fidelities. The average relative rank plots are computed over the validation error of the incumbent at  $z_{\max}$ , as obtained from the benchmark. For each seed, we obtain the ranks of algorithms for each benchmark, averaging them to get an estimate of each algorithm’s robustness across tasks. By computing the mean and standard error across seeds, we obtain an expectation of the algorithm’s performance with respect to its stochastic components and other variations derived from seeding. We show all relative ranks only after the  $1\times$  budget has been exhausted.

We compute ranks over the error or loss on the validation sets, where the choice of benchmark determines the exact metric to be minimized. Given we aggregate over the average ranks on a benchmark, the choice of per benchmark metric is irrelevant. For all prior-based baselines, we evaluate the mode of the prior distribution at  $z_{\max}$  as the first evaluation.

## E Algorithm details

### E.1 Pseudocode

Algorithm 3 is the main HB loop and the optimizer interface for PriorBand. Depending on the actual framework where this is implemented, Lines 8-9 may be arranged differently. L8 is an asynchronous call that allows HB and by extensions PriorBand to be parallel since different SH brackets can be scheduled simultaneously. That is, it can begin a new SH bracket even if L8 has not returned the sampled and evaluated base rung of the current SH bracket. L4 in Algorithm 4 can also be an asynchronous call, allowing multiple configurations from L3 to be evaluated in parallel. The pseudo-code for HB (Algorithm 3) thus presents itself in a way where the ensemble sampling policy  $\mathcal{E}_\pi$  replaces vanilla random search, for the lowest fidelity (base rung) in the current SH bracket in HB (L3 in Algorithm 4). Once the samples are collected and evaluated, vanilla-SH’s promotion can be run normally to survive and continue training chosen configurations.

---

**Algorithm 3** HB base for PriorBand

---

```
1: Input: Distribution over optimum  $\pi(\cdot)$ , halving parameter  $\eta$ , resource bounds  $[z_{\min}, z_{\max}]$ .
2:  $s_{\max} \leftarrow \lfloor \log_\eta(z_{\max}/z_{\min}) \rfloor$ 
3:  $\mathcal{H} \leftarrow \emptyset$ 
4: while budget remains do
5:   for  $s \in \{s_{\max}, \dots, 0\}$  do
6:      $n \leftarrow \left\lceil \frac{(s_{\max}+1)}{s+1} \cdot \eta^s \right\rceil$ ,  $\mathbf{z} \leftarrow z_{\max} \cdot \eta^{-s}$ 
7:      $r = s_{\max} - s$ 
8:      $\mathcal{H} \leftarrow \text{Alg. 4}(s, s_{\max}, \eta, n, \mathbf{z}, \pi, \mathcal{H})$  ▷ sample and evaluate  $n$  times
9:      $\mathcal{H} \leftarrow \text{Alg. 5}(s, s_{\max}, \eta, \mathbf{z}, \mathcal{H})$  ▷ run SH-based promotions
10:   end for
11: end while
```

---

---

**Algorithm 4** Sample and evaluate  $n$  configurations

---

```
1: Input:  $s, s_{\max}, \eta$ , number of samples  $n$ , evaluation resource  $z$ , prior  $\pi(\cdot)$ , observations  $\mathcal{H}$ 
2: for  $i \in \{1, 2, \dots, n\}$  do
3:    $\lambda \leftarrow \text{Alg. 1}(s, s_{\max}, \eta, \mathcal{H}, \pi(\cdot))$  ▷ sampling from  $\mathcal{U}(\cdot)$  here runs vanilla-HB
4:    $y \leftarrow \text{evaluate}(\lambda, z)$ 
5:    $\mathcal{S} \leftarrow \mathcal{S} \cup \{(\lambda, y, z)\}$ 
6: end for
7: return  $\mathcal{S}$ 
```

---

---

**Algorithm 5** Perform a SH iteration given the base rung

---

```
1: Input:  $s, s_{\max}, \eta$ , minimum evaluation resource  $z$ ,  $\mathcal{H}$ 
2:  $\text{rung} = s_{\max} - s$ 
3:  $(\hat{X}, \hat{Y}) \leftarrow$  retrieve all observations from  $r$  in  $\mathcal{H}$ 
4: for  $\hat{s} \in \{s - 1, \dots, 0\}$  do
5:    $z = \eta \cdot z$  ▷ the next higher fidelity
6:    $\hat{X} \leftarrow \text{top}_1/\eta(\hat{X}, \hat{Y})$  ▷ the top-performing  $(1/\eta)$  configurations in  $\hat{X}$ 
7:    $\hat{y} \leftarrow \text{evaluate}(\lambda, z), \forall \lambda \in \hat{X}$ 
8:    $\mathcal{S} \leftarrow \mathcal{S} \cup \{(\hat{X}, \hat{Y}, z)\}$ 
9: end for
10: return  $\mathcal{S}$ 
```

---

## E.2 Ablations

This section contains ablations over the various design choices in PriorBand (Section 5.2). All ablations are compared over different qualities of prior strength. Additionally, the set of benchmarks is also grouped into combinations of high-low correlation of lower fidelity to the  $z_{\max}$  and good-bad priors as discussed in Appendix D.1.5 and D.3. In Sections E.2.1-E.2.3 the next 3 sections we show ablations over different choices for the main design components in ESP for PriorBand. In Appendix E.2.4 we show the benefits of incumbent sampling in PriorBand.

### E.2.1 Random sampling proportions

Section 5.2.1 describes how the proportion of random samples is traded-off with the proportion of prior sampling across fidelities. In Figure 14 we compare 2 other designs for how the proportion of random and prior samples can be determined. PriorBand(50%) uses a fixed heuristic where 50% of the samples at each rung is going to be a random sample. PriorBand(linear) is similar to PriorBand, where instead of a geometric decay (Section 5.2.1) of random sample proportions like the latter, a linear decay is applied. That is, at the lowest rung,  $p_{\pi} = p_{\mathcal{U}}$ , and at the highest rung,  $p_{\pi} = \eta \cdot p_{\mathcal{U}}$ . For all intermediate rungs, the relationship is derived from this interpolated line.

### E.2.2 Prior and incumbent-based sampling trade-off

The choice of incumbent-prior trade-off strategy (Section 5.2.2) is crucial to the robustness of PriorBand. As the simplest design, we kept the proportion of incumbent sampling to prior sampling a fixed constant in PriorBand(constant),  $p_{\pi} = \eta \cdot p_{\hat{\lambda}}$ . PriorBand(decay)<sup>8</sup> in Figure 15 was designed such that the prior is decayed over time, similar to  $\pi$ BO. After every SH bracket, the proportion of prior samples with respect to the incumbent proportion was halved. As Figure 15 highlights, the data-driven likelihood score-based method described in Section 5.2.2 is essential for robust performance across all cases. This further supports our justification for the need for a non-naive method to achieve robustness, in Section 4.

### E.2.3 Choice of incumbent-based sampling

In this section, we look at different methods of defining an incumbent-based sampler. The role of the incumbent sampler is to allow exploitation by performing a local search around the best-seen

---

<sup>8</sup>this PriorBand(decay) is different from the decay version in Figure 14 in Appendix E.2.1

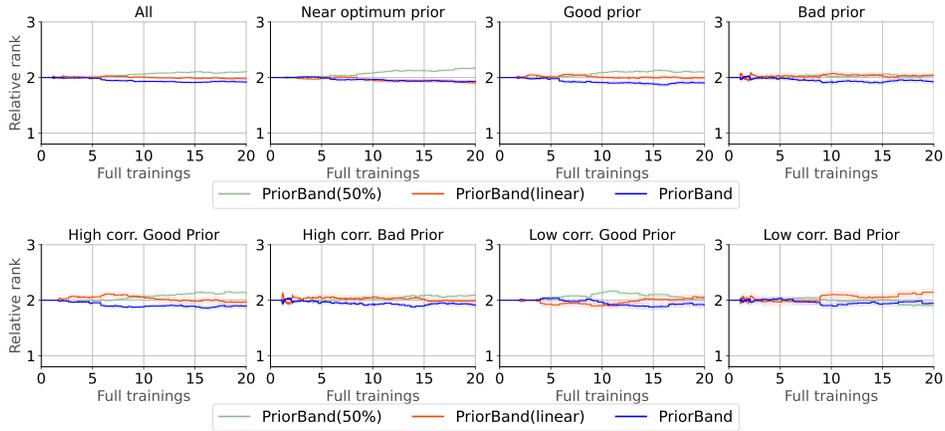


Figure 14: A comparison of `PriorBand` with 3 different strategies for setting the proportion of random samples in ESP as a fixed function of the rung. `PriorBand(50%)` fails to meaningfully utilize the *near optimum* and *good* prior when compared to both linear and geometric decay, with no aggregated benefit seen at any budget. `PriorBand(linear)` is an equivalent heuristic to `PriorBand` in intuition and simplicity, but overall `PriorBand`'s geometric decay comparatively performs robustly across all scenarios.

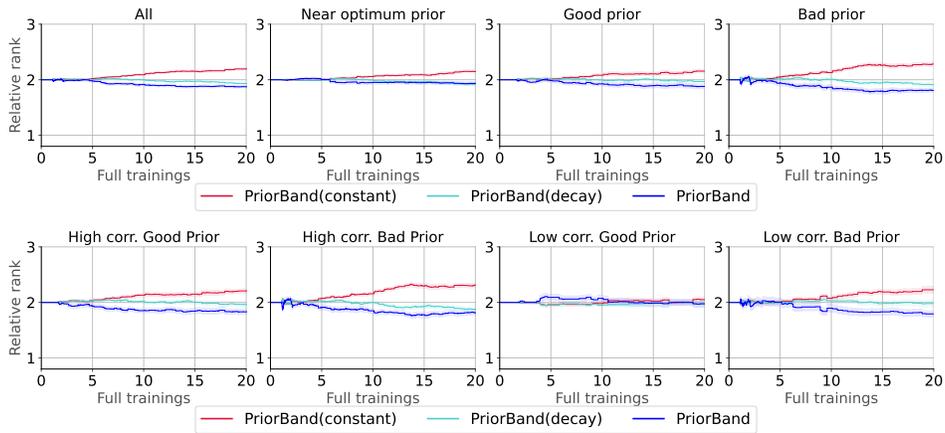


Figure 15: A comparison of `PriorBand` with 3 different strategies for trading off between incumbent sampling and prior distribution sampling. `PriorBand`'s default of using likelihood scores for weighting shows a dominant performance in almost all scenarios. The one exception is in low correlation settings with a good prior where all versions seem to suffer under low correlations. However, under the same set of tasks, a bad prior induces a marked difference in performance. In general, the plot highlights that each variant reacts differently to bad priors. The adaptive nature of `PriorBand` clearly is robust to different scenarios comparatively.

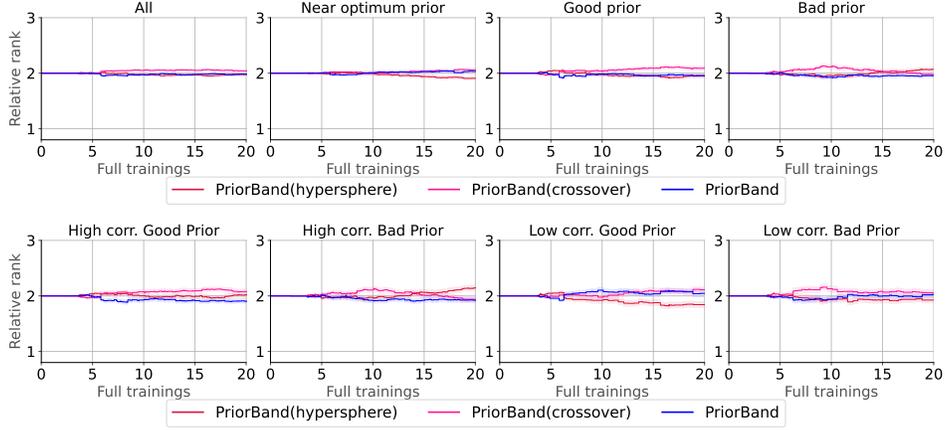


Figure 16: A comparison of 3 different methods in `PriorBand` to sample based on the incumbent configuration. No single method dominates in terms of performance. The hypersphere sampling method shows the most variation across scenarios which is undesirable for robustness while uniform crossover requires a less intuitive and complex operation. `PriorBand` uses local perturbation which is more intuitive, a much simpler operation, and shows similar performance to the crossover variant.

configuration. `PriorBand(hypersphere)` does so by sampling uniformly from a hypersphere around the incumbent, of a radius equivalent to the distance of the incumbent to its nearest neighbor. This method is intuitive, however, in practice, numerical issues may arise if the incumbent and its nearest neighbor are close to each other. This can make the sampling procedure extremely slow as the radius will keep shrinking, the more local search is performed. Moreover, the distance measure for different hyperparameter types can be an extra design choice. `PriorBand(crossover)` leverages the likelihood scores computed in Section 5.2.2 to choose if a random sample or a sample from the prior should participate in a simple uniform crossover with the incumbent. Though this alleviates the computation issue of the *hypersphere* method, it adds more complexity, is more difficult to interpret, and may not be as exploitative as desired of local search. Our final choice described in Section 5.1 is not only simple, fast, and easy to implement but also shows competitive performance with the other methods, as shown in Figure 16.

We additionally ablate over the chosen local mutation operation hyperparameters in Figure 17 and Figure 18.

#### E.2.4 Importance of incumbent-based sampling

In order to demonstrate the role of incumbent-based sampling in `PriorBand`'s design, Figure 19 shows a version of `PriorBand` where  $p_{\hat{\lambda}}$  is set to 0 with  $p_{\pi}$  inheriting all its probability (`PriorBand(w/o inc)`). Thereby, switching off the incumbent sampling, while keeping all other designs intact. We see that for the chosen geometric trade-off between prior and random sampling (Section 5.2.1, Appendix E.2.1), the inclusion of incumbent-based local search is essential.

#### E.2.5 Hyperparameters of incumbent-based sampling

For the incumbent-based sampling procedure described in Section 5.1, we employ a local perturbation around the incumbent using sampling from a Gaussian,  $\mathcal{N}(\hat{\lambda}, \sigma^2)$ . Here, we assume a fixed standard deviation (with  $\sigma = 0.25$ ) for the distribution. However, to make the perturbation more local, especially in higher dimensional spaces, we randomly choose if an HP will be perturbed or not, with a fixed probability of  $p = 0.5$ , an unbiased coin toss.

For the discrete hyperparameters, given  $k$  categorical choices  $C = \{c_1, c_2, \dots, c_k\}$ , we use a discrete distribution to perturb the category in the incumbent configuration. We give a weight of  $k$  to the current incumbent's categorical choice  $c_j$  and weight of 1 to each other choice, giving us sampling probabilities of  $p(c_j) = \frac{k}{2k-1}$  and  $p(c_i) = \frac{1}{2k-1}, \forall c_i \in C, c_i \neq c_j$ . This sampling procedure is hyperparameter free.

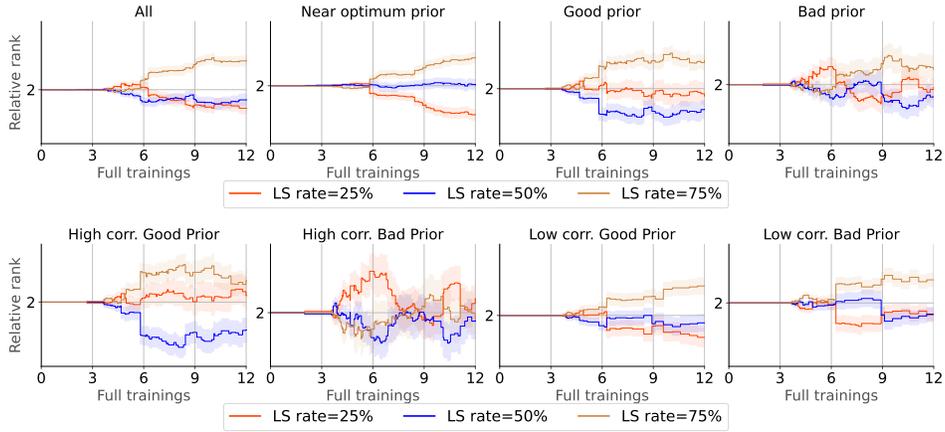


Figure 17: An ablation study over the hyperparameter  $p$  in the incumbent-based sampling in PriorBand which chooses the probability of selection of an HP to be perturbed for search. PriorBand chose  $p = 0.5$  which corresponds to 50% in the figure. This setting appears to be never the worst method. A local search (LS) rate of 25% is exploitative and can yield better performance under certain scenarios. However, our choice of 50% was chosen to be generally balanced. A high rate of 75% is too explorative and does not thus utilize the gains that the incumbent-based search can provide.

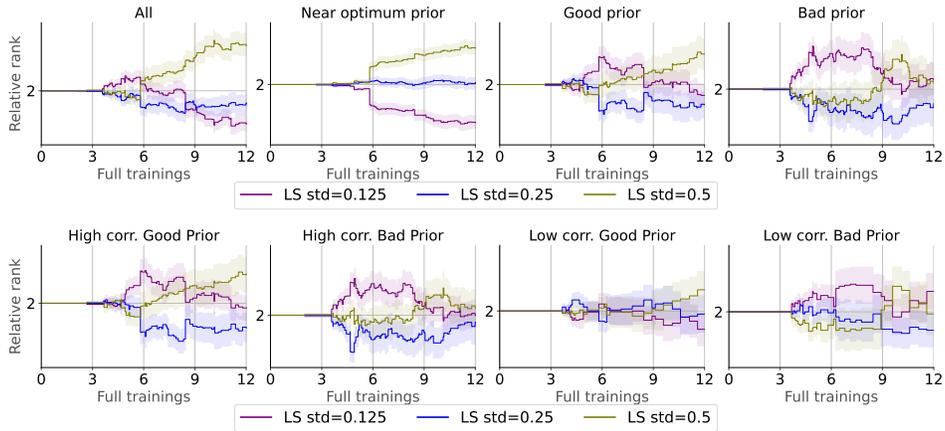


Figure 18: An ablation study over the hyperparameter  $\sigma$  in the incumbent-based sampling in PriorBand which chooses the standard deviation of the Gaussian that will be centered around the HP to be perturbed for local search (LS). PriorBand chose  $\sigma = 0.25$ . This setting appears to be never the worst method. As expected, a peaker distribution under  $\sigma = 0.125$  leads to more exploitation and thus strong performance under near-optimum priors. However, the quality of an incumbent improves over time. Being too exploitative, too early, could hurt optimization which is seen under the varying quality of good priors. The broad distribution under  $\sigma = 0.5$  is much more explorative and is thus relatively worst under good priors and competitive under bad priors.

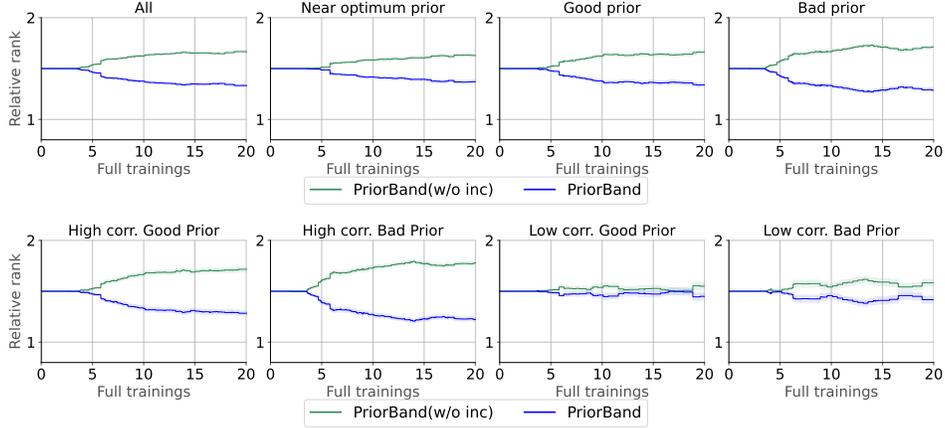


Figure 19: The top row shows different qualities of priors with the leftmost subplot being their aggregation. The bottom row shows different combinations of low and high correlation across fidelities in both the good and bad prior setting. The plot evidently shows the strength of incumbent-based sampling in all presented scenarios. The top row demonstrates how incumbent-based sampling helps escape the prior, which has less effect in the near optimum setting but a much stronger effect in the bad prior setting, where incumbent-based sampling is essential. In the bottom row, we see that incumbent-based sampling is more important in high correlation settings, where incumbents in low fidelities are likely to carry to later fidelities. The low correlation setting contrasts this effect, where spending the budget on configurations near the incumbent at low fidelities is not likely to translate to high-performing configurations at the highest fidelity.

The goal of `PriorBand` is to adapt to use good priors but recover from bad priors. To establish this through different experiments on different task settings and scenarios, we reduce confounding factors, keeping the aforementioned design for local search fixed.

Instead, hyperparameters could dynamically be adjusted to control `PriorBand`'s behavior, for example:

- A schedule to reduce the standard deviation of the Gaussian to sample the perturbation noise.
- Dynamically adjusting the probability  $p$  for HP perturbation as optimization proceeds, potentially yielding better exploitation.
- Expert prior insight into HPO landscapes can allow more custom setting of the standard deviation for the Gaussian distribution.
- Expert knowledge of interaction effect or the importance of the hyperparameters in the search space could also allow for a tuned setting of  $p$  for the selection of hyperparameters to perturb.

There are many possibilities for design, however, as our goal was to design the simplest approach that fulfills our desiderata, we choose a reasonable fixed setting and show it to be a robust choice in our experiments.

### E.2.6 Role of sampling the prior mode

Comparing `PriorBand`, where the prior mode is sampled at the maximum fidelity, with a version of `PriorBand` where the prior mode is sampled at the minimum fidelity (`Mode@min`), and not sampling the mode at all (`No-Mode`), in Figure 20. It appears that sampling the mode, in the beginning, provides huge gains especially if the prior is of good quality. In the case when the prior is not informative, `PriorBand` shows it can recover well even then. Given that most practical settings have multiple workers available, the fact that `PriorBand` can recover rapidly from misleading priors, and ultimately the utility of a good expert prior, it is reasonable to simply evaluate the expert default as the first evaluation. Under a multi-worker setup, the cost of this evaluation is amortized with the cheaper evaluations proceeding with the other workers. Optionally, the choice of whether the prior mode

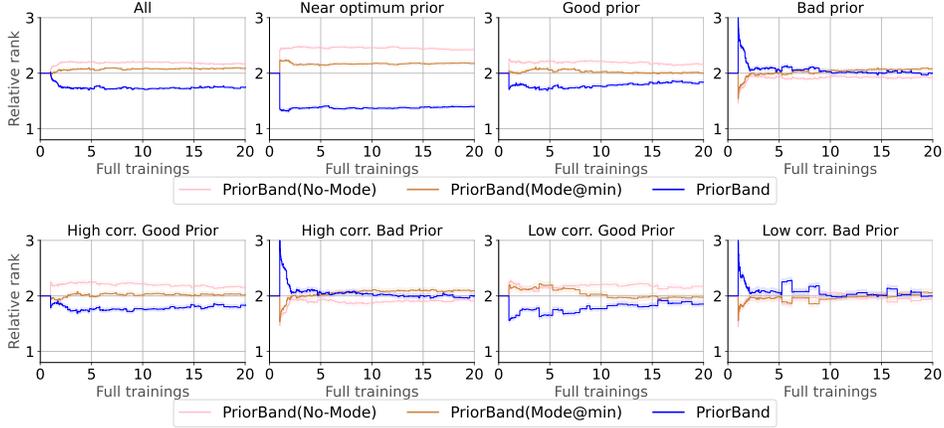


Figure 20: We show 3 possible methods of how to take the first evaluation. `PriorBand` by default choosing to sample the prior’s mode at the maximum fidelity, with `(Mode@min)` doing so at the minimum fidelity and `(No-Mode)` simply beginning with a random sample at the lowest fidelity. The most prominent failure case of `PriorBand` happens in the bad prior settings, as no exploration occurs until after 1 full training worth of budget is exhausted. However, the algorithm still recovers rapidly.

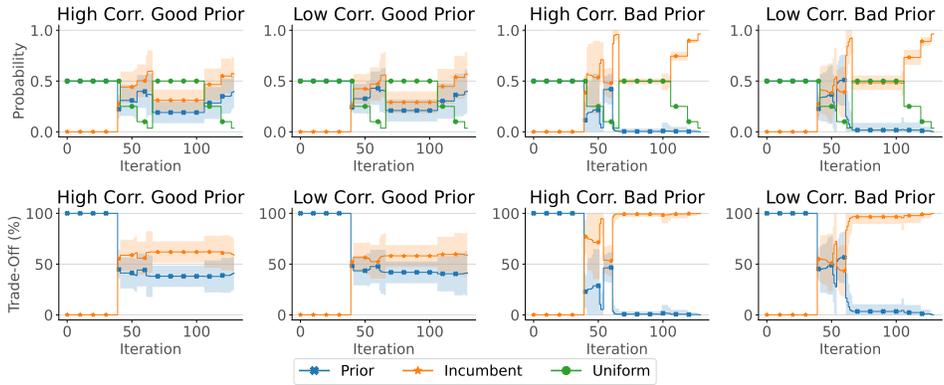


Figure 21: Evolution of the Ensemble Sampling policy  $\mathcal{E}_\pi$  and its probabilities  $p_U$ ,  $p_\pi$  and  $p_{\hat{\lambda}}$  for runs on the multi-fidelity Hartmann-3 benchmark, for all combinations of high/low fidelity correlation and good/bad prior strength. The x-axis counts the number of multi-fidelity evaluations made by HB. **[Top]** The y-axis shows the exact probability assigned to each sampling strategy,  $p_U$ ,  $p_\pi$ , and  $p_{\hat{\lambda}}$ . Uniform sampling follows a repeating pattern as a function of fidelity. Incumbent sampling remains inactive at the beginning, but once activated, it dynamically adjusts depending on the quality of the prior given. **[Bottom]** We visualize the dynamic trade-off between prior and incumbent-based sampling, showing  $p_{\hat{\lambda}}$  and  $p_\pi$  as a percentage of  $p_{\hat{\lambda}} + p_\pi$  on the y-axis. Under good priors, the probability of sampling from around the incumbent and sampling from priors is almost similar. Whereas under bad priors, the prior is discarded almost completely after one complete round of HB. Differences across high-low correlation setups exist, in the period between activation of the incumbent and one complete iteration of HB, but values stabilize after this iteration for all cases.

should be evaluated as the first evaluation or not can be toggled. For example, when the expert is confident of a good prior configuration and has knowledge of its performance from previous runs.

### E.3 Post-hoc view of `PriorBand` for interpretability

A DL expert can often provide a prior  $\pi(\cdot)$  but is often unsure as to its merits. By tracking the evolution of `PriorBand`’s sampling probabilities during the optimization process, we can get an indicator of the strength of the  $\pi(\cdot)$ . The more useful the prior is for performance, the higher the probability of sampling from  $\pi(\cdot)$  and thus a higher  $p_\pi$ . We illustrate this in Figure 21, where we plot

the sampling probabilities  $p_{\mathcal{U}}$ ,  $p_{\pi}$  and  $p_{\hat{\lambda}}$  of `PriorBand` during an HPO run with 50 seeds on the synthetic multi-fidelity Hartmann-3 function (Appendix D.1.1). We clearly see that our motivation for the design choice of `PriorBand` is well supported as the probability of sampling from the incumbent is not suppressing sampling from prior under good priors, but are aggressively affecting the chance of sampling from bad priors. Such a post-hoc analysis offers a DL expert insights into their own prior knowledge, allowing them to update or re-enforce their beliefs about what a good prior is for the problem at hand.

#### E.4 Model extensions

In this section, we elucidate the model-based components when extending algorithms with  $\mathcal{E}_{\pi}$  as shown in Figure 7. Firstly, we briefly explain BOHB and its modeling choice. BOHB subsumes the hyperparameter of the initial design size by setting it to  $N_{dim} + 2$  where  $N_{dim}$  is the dimensionality of the search space of a task. To activate model-based search, BOHB uses the following criteria: find the highest fidelity level with at least  $N_{dim} + 2$  evaluated observations. If no such fidelity level exists, continue with uniform random sampling. If such fidelity exists, use all the observations at that fidelity to build a TPE as the surrogate. Since a model is built at a fidelity level, the fidelity variable is not part of the feature set for the surrogate. During acquisition, EI is used to obtain a configuration that approximately maximizes the TPE surrogate.

$$\text{EI}_z(\mathbf{x}|\mathcal{D}) = \mathbb{E}[\max\{f_z(\mathbf{x}) - y_z^{min}, 0\}], \quad (6)$$

where  $y_z^{min}$  is the best score seen at the fidelity  $z$ . Given  $f$  is a model built at a fidelity level  $z$ , the EI acquisition estimates the improvement of the suggested configuration at fidelity  $z$ . When the number of observations at  $z_{max}$  is  $N_{dim} + 2$ , the EI acquisition performs similarly to vanilla-BO. Mobster or asynchronous-BOHB follows the exact BO loop as BOHB, except that it uses a GP instead of a TPE and uses asynchronous HB for scheduling and not vanilla-HB.

**Model extension to `PriorBand`.** In our experiments to extend `PriorBand` with a model in `PriorBand+BO`, the above procedure of automatically switching to model-search lead to  $\mathcal{E}_{\pi}$  not taking action and affecting optimization as we desire.

Similar to the initial design size in BO, incumbent-based sampling has an ‘activation criteria’ of one full SH bracket being evaluated, and at least one configuration evaluated at  $z_{max}$ , after which incumbent-based search begins. We could default to BOHB’s ‘activation criteria’ but for certain problems, the number of observations would satisfy the  $N_{dim} + 2$  criteria even before the 1<sup>st</sup> SH bracket is over. This implies that incumbent-based sampling, a crucial component of  $\mathcal{E}_{\pi}$ , is never activated. Hence, we follow an approach that is more in line with BO and  $\pi$ BO in which 10 function evaluations are used as the initial design budget before switching to model search. This is most evident in Figure 7, where  $\pi$ BO behaves identically to only performing Random Search on the prior distribution, diverging at 10 full function evaluations.

We treat  $10 \times$  as  $(10 \cdot z_{max})$  the total budget (in epochs) exhausted during multi-fidelity optimization. After which, a GP model is activated for search that models *all* the observations made during the optimization, across any fidelity available. That is, the fidelity is an extra dimension modeled along with the search space. During the acquisition, since it is known from the optimization state  $s_t$  which fidelity  $z$  the new sample will be evaluated at, a 2-step optimization is performed when maximizing EI. In the first step, a set of configurations (10 in our experiments) is extracted for fidelity  $z$  through Monte Carlo estimates of Equation 6 returning configurations likely to improve over the best configuration found so far at  $z$ . Following this, the EI score is calculated for this selected set of configurations using Equation 6 but with  $z = z_{max}$ . At this stage,  $y^{min}$  is chosen to be the best score obtained across all observations. The idea is to choose a configuration that is likely to maximize performance at the fidelity level where it is being queried and is also likely to improve at the target fidelity  $z_{max}$ .

**Contextualizing model-search under  $\mathcal{E}_{\pi}$ .** During the typical initial design phase of `PriorBand+BO`, it is `PriorBand` that runs with the defined  $\mathcal{E}_{\pi}$  comprising of  $\mathcal{U}(\cdot)$ ,  $\pi(\cdot)$ ,  $\hat{\lambda}(\cdot)$  as actions. At the end of the initial design phase of `PriorBand+BO`, the action set  $\mathcal{A}$  updates from  $\{\pi(\cdot), \mathcal{U}(\cdot), \hat{\lambda}(\cdot)\}$  to  $\{\pi(\cdot), \mathcal{U}(\cdot), \hat{\lambda}(\cdot), \mathcal{M}(\cdot)\}$  with new weights as  $p_{\pi} = p_{\mathcal{U}} = p_{\hat{\lambda}} = 0$  and  $p_{\mathcal{M}} = 1$ , where  $\mathcal{M}$  denotes model-based search.

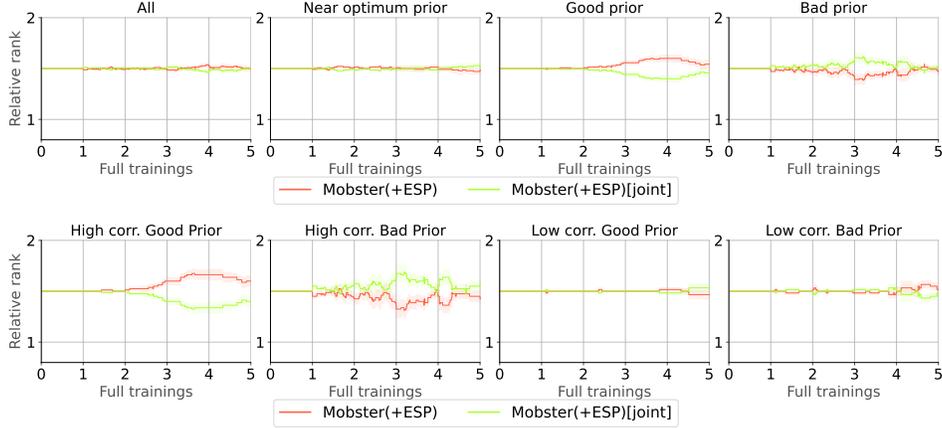


Figure 22: Comparing multi-fidelity modelling on asynchronous-HB per *rung* and *jointly* over the search space and fidelity (j). (Top) Compares the performances over different qualities of priors. (Bottom) Compares the good (at25) and bad prior cases, under benchmarks grouped into good-bad correlation. All algorithms were run for a total of  $20\times$  over 4 workers.

**Comparing multi-fidelity modeling.** In the previous section we motivate criterion and acquisition that allows  $\mathcal{E}_\pi$  to influence search and the initial design space. For `PriorBand`, we note that the incumbent activation criteria are not fulfilled if modeling per fidelity with dimensionality as a criterion, as is done in `BOHB` and `Mobster`. In contrast, asynchronous-HB samples at random fidelities instead of the lowest fidelity first. When used with the ESP this allows the possibility of incumbent-based samples being activated before the model search begins. Hence, for asynchronous HB we can apply per-fidelity modeling. In Figure 22, we thus compare the 2 types of modeling discussed above: per-fidelity (`Mobster+E`) and joint modeling like `PriorBand+BO` (`Mobster+E[joint]`). We conclude that there is no significant difference in the performance of the two modeling choices. Though the joint model seems to perform slightly better under good priors.

## F More on experiments

This section gives a detailed experiment analysis that expounds on the results from Section 7.1-7.3. We compare similar setups but include our constructed prior-based baselines, over a different grouping of benchmarks with high and low correlations (Appendix F.1, F.2, F.3). In Appendix F.4 we compare the prior-based algorithms with each other. We further show the final validation performance tables of algorithms for the different experiments for every benchmark-prior combination, over 2 budget horizons in Appendix F.5.

### F.1 Robustness of `PriorBand`

More supporting results for Figure 5 in Section 7.1. we show the same results but with the addition of *near optimum* priors. Figure 23 shows the relative rank comparison over near optimum, good, and bad priors, as well as the aggregate (*All*). Most notably, we see the relation between `PriorBand` and HB with respect to prior strength, showcasing substantial benefits of strong priors while recovering in the bad prior setting.

To further illustrate how different tasks affect optimizer behavior, we cluster the set of 12 benchmarks we’ve chosen into 8 good correlation benchmarks and 4 bad correlation benchmarks as defined in Appendix D.1.5. Figure 23 shows the relative ranks for the same set of algorithms when grouped and aggregated along these criteria. `PriorBand` is the most robust algorithm shown here.

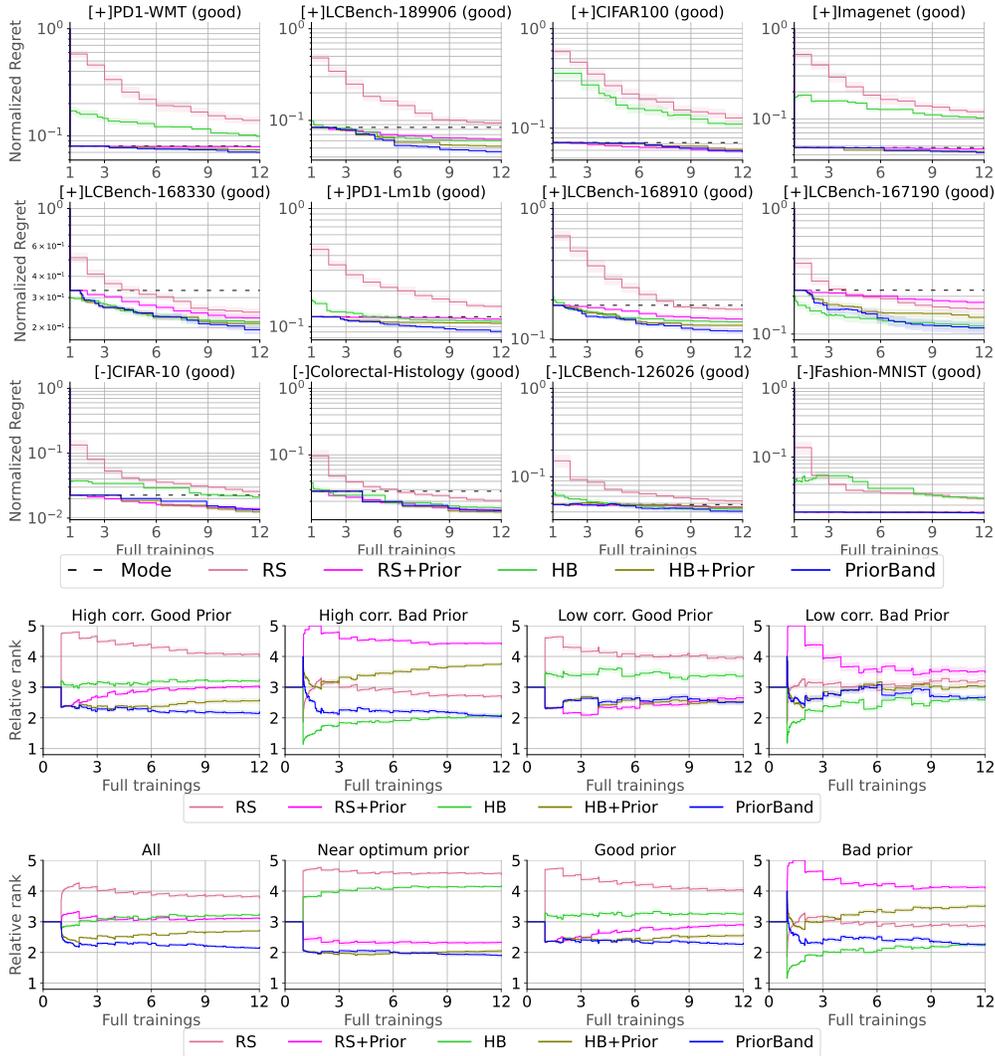


Figure 23: [Top] Each plot compares the algorithms based on the average normalized regret across 50 seeds, under the good prior setting. The optima for a benchmark was assumed to be the best scores achieved by all algorithms across all seeds. PriorBand is anytime best across in all cases. [+ ] denotes the benchmark tends to have a strong correlation across fidelities, [- ] denotes a weak or poor correlation across fidelities; [Bottom] Comparing PriorBand with other baselines on single-worker runs for 50 seeds. The *top* row compares the average relative ranks across all benchmarks under different prior qualities, where the *All* averages over the 3 prior strengths too. The *bottom* row groups the benchmarks into high-low correlation sets based on fidelity correlations and creates 4 scenarios when combined with good-bad priors. PriorBand emerges as the most consistent performer across all 7 scenarios. Every other baseline has at least one setting where it is one of the 2 worst algorithms. Prior-based methods show improved performance with a good prior but suffer with a bad prior, as expected. Likewise, multi-fidelity methods benefit from high correlation. Under low-correlation settings, the ranking gap between RS and HB expectedly is lower.

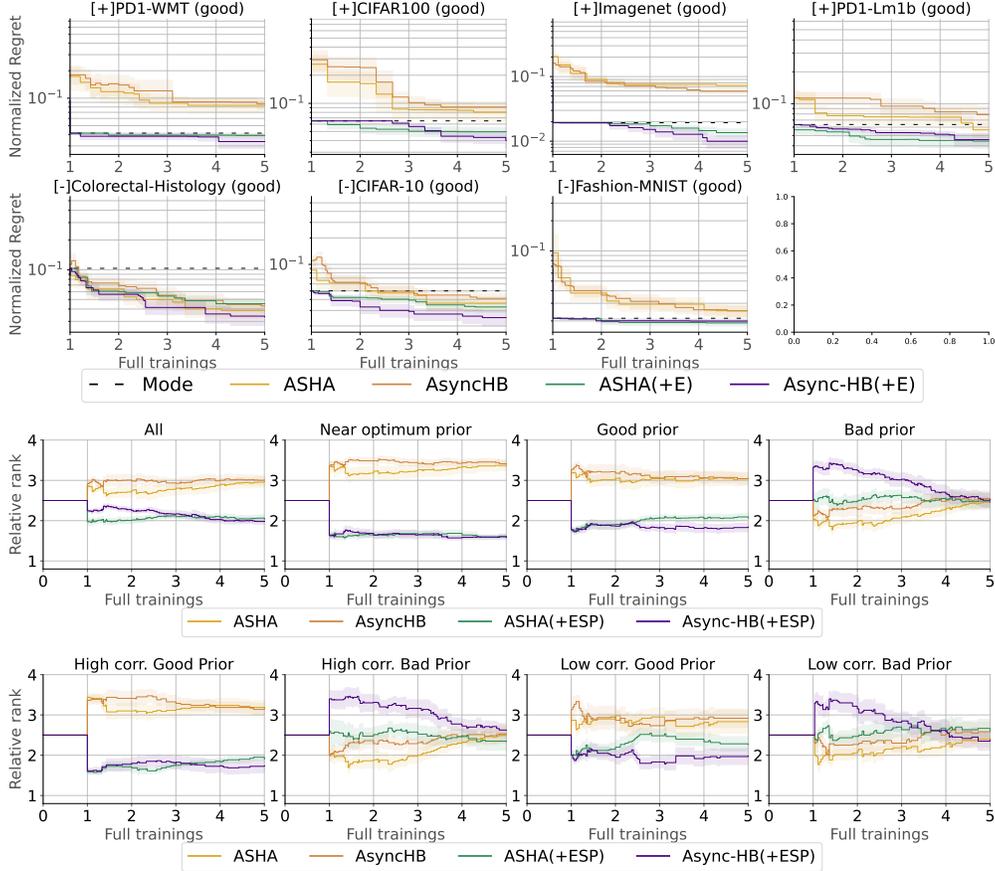


Figure 24: [Top] Each plot compares the algorithms based on the average normalized regret across 50 seeds, under the good prior setting, run over 4 workers each. The optima for a benchmark was assumed to be the best scores achieved by all algorithms across all seeds. [+] denotes the benchmark tends to have a strong correlation across fidelities, [-] denotes a weak or poor correlation across fidelities. Refer to Appendix D.1.6 for missing benchmark; [Bottom] Comparing ASHA, AsynchronousHB with their ESP augmented versions (+E) and PriorBand when distributed over 4 workers for a total budget of  $20\times$ . The *top* row compares the average relative ranks across all benchmarks under different prior qualities, where the *All* averages over the 3 prior strengths. The *bottom* row groups the benchmarks into high-low correlation sets based on fidelity correlations and creates 4 scenarios when combined with good-bad priors. All 3 algorithms with the ESP  $\mathcal{E}_\pi$  show superior performance under good priors and strong recovery under bad priors. Under a bad prior for good correlation benchmarks, the recovery of ESP-based asynchronous methods is slower since the vanilla algorithms perform better under good correlation. Unlike the low correlation set under bad priors, where all  $\mathcal{E}_\pi$ -based algorithms show faster recovery. Overall, PriorBand remains a strong performer even in the parallel setup.

## F.2 Generality of $\mathcal{E}_\pi$

Figure 24 (top) compares 3 different prior qualities for more supporting results for Figure 6 in Section 7.2. Figure 24 (bottom) splits the set of 12 benchmarks into high-low correlations and plots their interaction with good-bad priors.

## F.3 Extensibility of $\mathcal{E}_\pi$ with models

More supporting results for Figure 7 in Section 7.3. In Figure 25, which compares 3 different prior qualities, we assess how the quality of the prior effects model-based methods. We also compare the

Table 10: Table comparing Random Search (RS), HyperBand (HB), and PriorBand’s final validation errors of the current incumbent at the highest fidelity at 2 budget horizons of  $5\times$  and  $12\times$ . Runs are averaged over 50 seeds on 1 worker each. PriorBand shows superior anytime performance under informative priors. Under extremely bad priors, for shorter compute budgets ( $5\times$ ) PriorBand has a poor start, however, given an adequate budget ( $12\times$ ) PriorBand can match HB’s performance on average.

Benchmark	5x			12x		
	RS	HB	PriorBand	RS	HB	PriorBand
<b>Near Optimum Prior</b>						
JAHS-C10	11.454 ± 1.935	11.129 ± 1.256	<b>8.252 ± 0.215</b>	10.451 ± 1.055	10.077 ± 0.769	<b>8.236 ± 0.197</b>
JAHS-CH	6.752 ± 0.943	6.332 ± 1.12	<b>4.449 ± 0.124</b>	5.977 ± 0.646	5.704 ± 0.659	<b>4.445 ± 0.127</b>
JAHS-FM	5.497 ± 0.396	5.846 ± 0.823	<b>4.724 ± 0.06</b>	5.282 ± 0.299	5.306 ± 0.267	<b>4.719 ± 0.055</b>
LC-126026	0.053 ± 0.011	0.046 ± 0.01	<b>0.024 ± 0.008</b>	0.048 ± 0.008	0.044 ± 0.008	<b>0.023 ± 0.006</b>
LC-167190	0.214 ± 0.021	0.193 ± 0.025	<b>0.136 ± 0.017</b>	0.2 ± 0.02	0.187 ± 0.023	<b>0.133 ± 0.013</b>
LC-168330	0.444 ± 0.046	0.41 ± 0.035	<b>0.29 ± 0.041</b>	0.406 ± 0.038	0.397 ± 0.032	<b>0.275 ± 0.028</b>
LC-168910	0.38 ± 0.093	0.314 ± 0.021	<b>0.2 ± 0.021</b>	0.324 ± 0.032	0.308 ± 0.018	<b>0.194 ± 0.019</b>
LC-189906	0.222 ± 0.09	0.151 ± 0.021	<b>0.113 ± 0.015</b>	0.166 ± 0.033	0.145 ± 0.019	<b>0.107 ± 0.011</b>
PD1-Cifar100	0.378 ± 0.156	0.338 ± 0.091	<b>0.257 ± 0.051</b>	0.302 ± 0.074	0.29 ± 0.039	<b>0.238 ± 0.04</b>
PD1-ImageNet	0.333 ± 0.108	0.306 ± 0.041	<b>0.239 ± 0.038</b>	0.282 ± 0.049	0.268 ± 0.022	<b>0.228 ± 0.026</b>
PD1-LM1B	0.681 ± 0.034	0.651 ± 0.013	<b>0.637 ± 0.015</b>	0.658 ± 0.013	0.648 ± 0.012	<b>0.632 ± 0.015</b>
PD1-WMT	0.454 ± 0.097	0.402 ± 0.032	<b>0.357 ± 0.029</b>	0.403 ± 0.038	0.384 ± 0.02	<b>0.347 ± 0.025</b>
<b>Good Prior</b>						
JAHS-C10	11.454 ± 1.935	11.129 ± 1.256	<b>9.986 ± 0.369</b>	10.451 ± 1.055	10.077 ± 0.769	<b>9.457 ± 0.415</b>
JAHS-CH	6.752 ± 0.943	6.332 ± 1.12	<b>5.974 ± 0.535</b>	5.977 ± 0.646	5.704 ± 0.659	<b>5.589 ± 0.413</b>
JAHS-FM	5.497 ± 0.396	5.846 ± 0.823	<b>5.037 ± 0.021</b>	5.282 ± 0.299	5.306 ± 0.267	<b>5.026 ± 0.044</b>
LC-126026	0.053 ± 0.011	0.046 ± 0.01	<b>0.045 ± 0.006</b>	0.048 ± 0.008	0.044 ± 0.008	<b>0.043 ± 0.007</b>
LC-167190	0.214 ± 0.021	<b>0.193 ± 0.025</b>	0.198 ± 0.026	0.2 ± 0.02	0.187 ± 0.023	<b>0.186 ± 0.026</b>
LC-168330	0.444 ± 0.046	<b>0.41 ± 0.035</b>	0.411 ± 0.024	0.406 ± 0.038	0.397 ± 0.032	<b>0.387 ± 0.035</b>
LC-168910	0.38 ± 0.093	0.314 ± 0.021	<b>0.313 ± 0.019</b>	0.324 ± 0.032	0.308 ± 0.018	<b>0.295 ± 0.023</b>
LC-189906	0.222 ± 0.09	0.151 ± 0.021	<b>0.147 ± 0.013</b>	0.166 ± 0.033	0.145 ± 0.019	<b>0.134 ± 0.013</b>
PD1-Cifar100	0.378 ± 0.156	0.338 ± 0.091	<b>0.258 ± 0.002</b>	0.302 ± 0.074	0.29 ± 0.039	<b>0.248 ± 0.013</b>
PD1-ImageNet	0.333 ± 0.108	0.306 ± 0.041	<b>0.224 ± 0.001</b>	0.282 ± 0.049	0.268 ± 0.022	<b>0.219 ± 0.008</b>
PD1-LM1B	0.681 ± 0.034	0.651 ± 0.013	<b>0.647 ± 0.008</b>	0.658 ± 0.013	0.648 ± 0.012	<b>0.642 ± 0.009</b>
PD1-WMT	0.454 ± 0.097	0.402 ± 0.032	<b>0.37 ± 0.008</b>	0.403 ± 0.038	0.384 ± 0.02	<b>0.367 ± 0.009</b>
<b>Bad Prior</b>						
JAHS-C10	11.454 ± 1.935	<b>11.129 ± 1.256</b>	11.729 ± 1.795	10.451 ± 1.055	10.077 ± 0.769	<b>10.075 ± 0.777</b>
JAHS-CH	6.752 ± 0.943	<b>6.332 ± 1.12</b>	6.553 ± 1.284	5.977 ± 0.646	5.704 ± 0.659	<b>5.551 ± 0.516</b>
JAHS-FM	<b>5.497 ± 0.396</b>	5.846 ± 0.823	5.962 ± 0.973	5.282 ± 0.299	5.306 ± 0.267	<b>5.277 ± 0.248</b>
LC-126026	0.053 ± 0.011	<b>0.046 ± 0.01</b>	0.051 ± 0.011	0.048 ± 0.008	<b>0.044 ± 0.008</b>	0.045 ± 0.01
LC-167190	0.214 ± 0.021	<b>0.193 ± 0.025</b>	0.196 ± 0.024	0.2 ± 0.02	0.187 ± 0.023	<b>0.183 ± 0.026</b>
LC-168330	0.444 ± 0.046	<b>0.41 ± 0.035</b>	0.424 ± 0.035	0.406 ± 0.038	<b>0.397 ± 0.032</b>	0.401 ± 0.042
LC-168910	0.38 ± 0.093	0.314 ± 0.021	<b>0.309 ± 0.019</b>	0.324 ± 0.032	0.308 ± 0.018	<b>0.297 ± 0.016</b>
LC-189906	0.222 ± 0.09	<b>0.151 ± 0.021</b>	0.167 ± 0.03	0.166 ± 0.033	<b>0.145 ± 0.019</b>	0.15 ± 0.027
PD1-Cifar100	0.378 ± 0.156	<b>0.338 ± 0.091</b>	0.473 ± 0.2	0.302 ± 0.074	<b>0.29 ± 0.039</b>	0.297 ± 0.058
PD1-ImageNet	0.333 ± 0.108	<b>0.306 ± 0.041</b>	0.315 ± 0.04	0.282 ± 0.049	<b>0.268 ± 0.022</b>	0.28 ± 0.03
PD1-LM1B	0.681 ± 0.034	<b>0.651 ± 0.013</b>	0.656 ± 0.014	0.658 ± 0.013	<b>0.648 ± 0.012</b>	0.648 ± 0.013
PD1-WMT	0.454 ± 0.097	<b>0.402 ± 0.032</b>	0.425 ± 0.057	0.403 ± 0.038	0.384 ± 0.02	<b>0.383 ± 0.021</b>

model-based extension of PriorBand (PriorBand+BO) against Mobster (asynchronous BOHB) in the distributed setting, in Figure 26.

#### F.4 Comparing prior-based baselines

We would like to determine if PriorBand is sufficient to outperform existing prior-based baselines as well as our own extension of PriorBand which is model-based, namely PriorBand+BO. Figure 27 show the results of this comparison.

#### F.5 Final performance tables

In this section, Tables 10, 11, 12, 13, 14, compare the performance of PriorBand, its BO extension and other applications of  $\mathcal{E}_\pi$ , to commonly used algorithms from the literature for the different classes of optimizers. The tables show the final performance at 2 budget horizons, highlighting the robustness of the ESP-based algorithms.

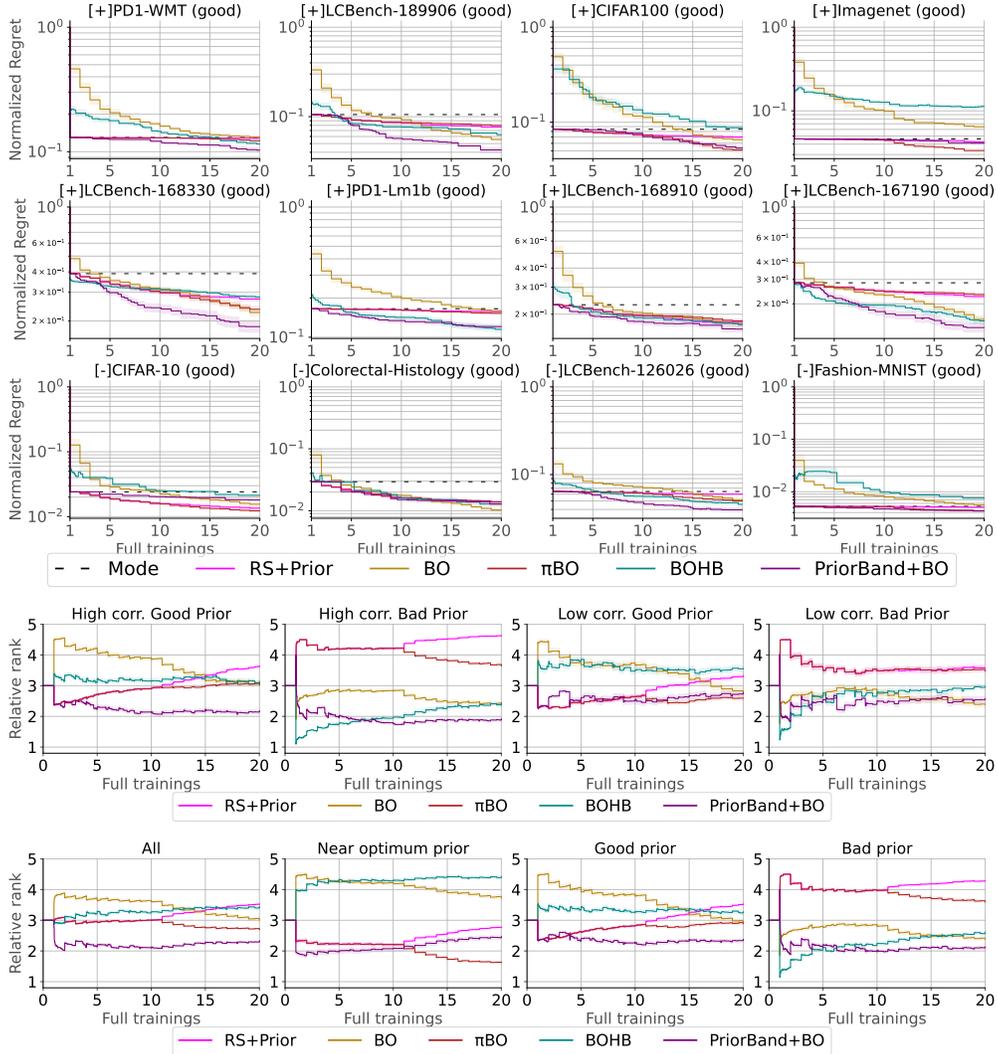


Figure 25: [Top] Each plot compares the algorithms based on the average normalized regret across 50 seeds, under the good prior setting. The optima for a benchmark was assumed to be the best scores achieved by all algorithms across all seeds. PriorBand+BO is anytime best across in all cases. [+] denotes the benchmark tends to have a strong correlation across fidelities, [-] denotes a weak or poor correlation across fidelities; [Bottom] Comparing model extensions over 3 sets of prior qualities. Using PriorBand with BO shows a dominant performance in almost all cases. In the first row, we show that  $\pi$ BO does outperform all other methods in the near optimum setting, once model sampling activates. This is likely due to  $\pi$ BO’s emphasis on the prior which is the same cause for its poor performance in the bad prior setting. Both BO and BOHB suffer from having no access to the prior in the near optimum/good prior setting while only BOHB really has an advantage over PriorBand+BO early on in the bad prior setting. In the second row, we see that extending Mobster with the ensembling policy  $\mathcal{E}_\pi$  imbues Mobster with an effective mechanism for exploiting priors. However, this exploitation has a more consistent negative impact when the prior is bad, at least when compared to vanilla Mobster and our model-based PriorBand+BO.

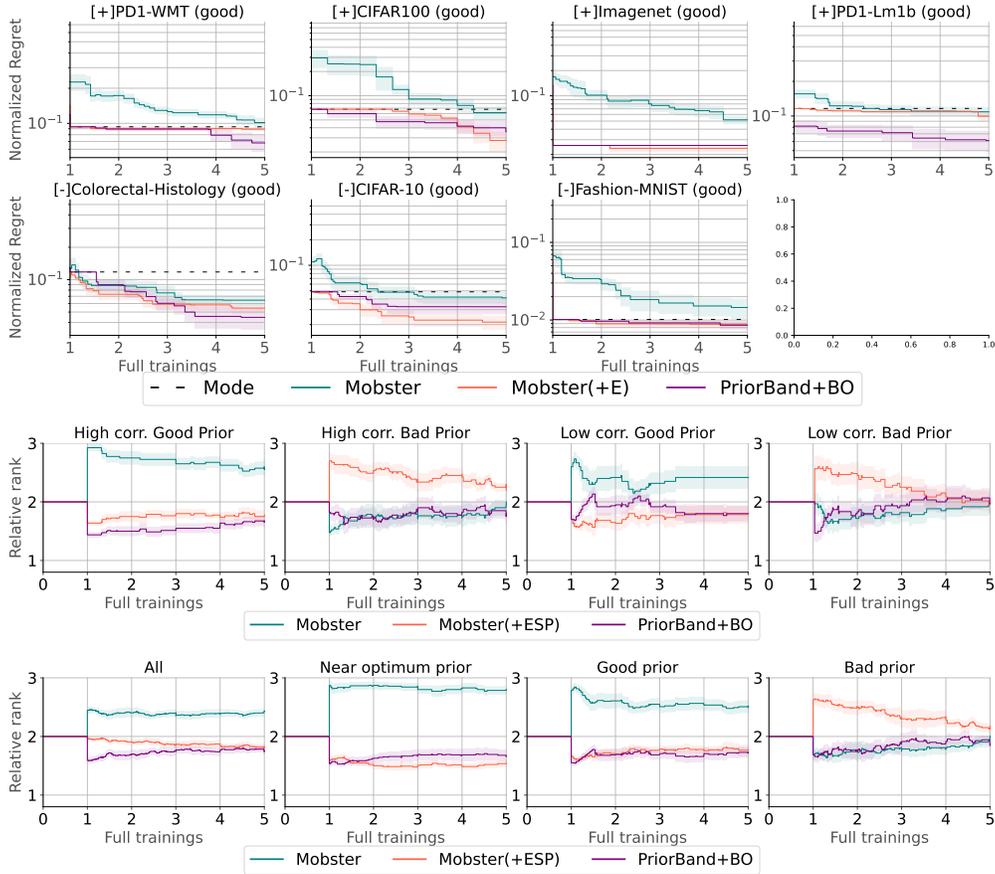


Figure 26: [Top] Each plot compares the algorithms based on the average normalized regret across 50 seeds, under the good prior setting, run over 4 workers each. The optima for a benchmark was assumed to be the best scores achieved by all algorithms across all seeds. [+] denotes the benchmark tends to have a strong correlation across fidelities, [-] denotes a weak or poor correlation across fidelities. Refer to Appendix D.1.6 for missing benchmark; [Bottom] Comparing model extensions over 3 sets of prior qualities. Using PriorBand with BO shows a dominant performance in almost all cases. In the first row, we show that  $\pi$ BO does outperform all other methods in the near optimum setting, once model sampling activates. This is likely due to  $\pi$ BO’s emphasis on the prior which is the same cause for its poor performance in the bad prior setting. Both BO and BOHB suffer from having no access to the prior in the near optimum/good prior setting while only BOHB really has an advantage over PriorBand+BO early on in the bad prior setting. In the second row, we see that extending Mobster with the ensembling policy  $\mathcal{E}_\pi$  imbues Mobster with an effective mechanism for exploiting priors. However, this exploitation has a more consistent negative impact when the prior is bad, at least when compared to vanilla Mobster and our model-based PriorBand+BO.

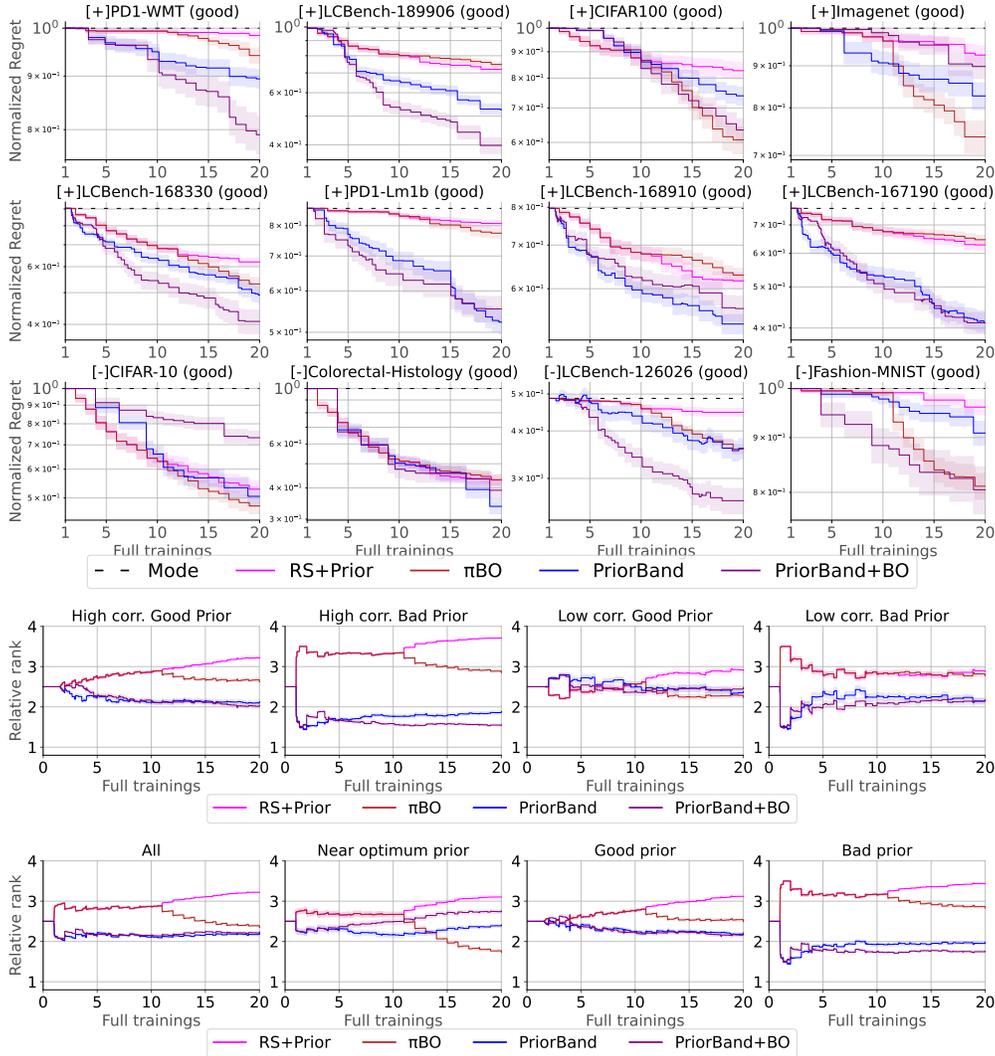


Figure 27: [Top] Each plot compares the algorithms based on the average normalized regret across 50 seeds, under the good prior setting. The optima for a benchmark was assumed to be the best scores achieved by all algorithms across all seeds. PriorBand+BO offers its model benefits over PriorBand under longer budgets but has mixed performance under poor correlation of fidelity performance. [+] denotes the benchmark tends to have a strong correlation across fidelities, [-] denotes a weak or poor correlation across fidelities; [Bottom] Comparing model-based methods which incorporate the prior. In the near optimum case,  $\pi$ BO’s strong performance only occurs once its model-based sampling occurs. This strong emphasis on searching around the prior is also what causes  $\pi$ BO’s weaker performance in the bad prior case. When comparing PriorBand against its model counterpart, PriorBand+BO, we see that PriorBand utilizes the prior better while PriorBand+BO with its model-based sampling leads to better performance in the long run if the prior is bad. This trade-off is averaged off in the case of a good prior which is neither overly optimistic as the near optimum prior is, nor neither as pessimistic as the bad prior is.

Table 11: Table comparing Asynchronous Successive Halving (ASHA), Asynchronous HyperBand (AsyncHB) and PriorBand final validation errors of the current incumbent at the highest fidelity at 2 budget horizons of  $1\times$  and  $5\times$ . Runs are averaged over 10 seeds where each run is with 4 workers. Unlike the other 2 algorithms, PriorBand does not performance ASHA-like asynchronous promotions to minimize idle workers. PriorBand simply starts the next SH bracket if a worker is free. The table shows that PriorBand can maintain its robust performance when run in a parallel setting while being competitive to related asynchronous algorithms.

Benchmark	$1\times$			$5\times$		
	Near Optimum Prior					
	ASHA	AsyncHB	PriorBand	ASHA	AsyncHB	PriorBand
JAHS-C10	$11.382 \pm 2.523$	$12.192 \pm 1.558$	<b><math>8.281 \pm 0.221</math></b>	$9.739 \pm 0.716$	$9.901 \pm 0.598$	<b><math>8.281 \pm 0.221</math></b>
JAHS-CH	$6.325 \pm 1.116$	$6.771 \pm 1.296$	<b><math>4.466 \pm 0.188</math></b>	$5.492 \pm 0.362$	$5.565 \pm 0.342$	<b><math>4.399 \pm 0.143</math></b>
JAHS-FM	$6.627 \pm 3.384$	$6.245 \pm 0.999$	<b><math>4.713 \pm 0.052</math></b>	$5.127 \pm 0.235$	$5.134 \pm 0.274$	<b><math>4.709 \pm 0.049</math></b>
PDI-Cifar100	$0.411 \pm 0.152$	$0.434 \pm 0.172$	<b><math>0.243 \pm 0.056</math></b>	$0.271 \pm 0.04$	$0.279 \pm 0.032$	<b><math>0.221 \pm 0.043</math></b>
PDI-ImageNet	$0.366 \pm 0.049$	$0.333 \pm 0.056$	<b><math>0.259 \pm 0.059</math></b>	$0.265 \pm 0.019$	$0.255 \pm 0.011$	<b><math>0.218 \pm 0.021</math></b>
PDI-LM1B	$0.665 \pm 0.016$	$0.663 \pm 0.015$	<b><math>0.643 \pm 0.027</math></b>	$0.649 \pm 0.008$	$0.655 \pm 0.016$	<b><math>0.628 \pm 0.008</math></b>
PDI-WMT	$0.446 \pm 0.09$	$0.443 \pm 0.06$	<b><math>0.362 \pm 0.033</math></b>	$0.396 \pm 0.026$	$0.402 \pm 0.027$	<b><math>0.349 \pm 0.025</math></b>
Good Prior						
	ASHA	AsyncHB	PriorBand	ASHA	AsyncHB	PriorBand
JAHS-C10	$11.382 \pm 2.523$	$12.192 \pm 1.558$	<b><math>10.194 \pm 0.0</math></b>	$9.739 \pm 0.716$	$9.901 \pm 0.598$	<b><math>9.413 \pm 0.42</math></b>
JAHS-CH	<b><math>6.325 \pm 1.116</math></b>	$6.771 \pm 1.296$	$6.603 \pm 0.0$	$5.492 \pm 0.362$	$5.565 \pm 0.342$	<b><math>5.238 \pm 0.325</math></b>
JAHS-FM	$6.627 \pm 3.384$	$6.184 \pm 1.009$	<b><math>5.042 \pm 0.0</math></b>	$5.127 \pm 0.235$	$5.126 \pm 0.211$	<b><math>5.024 \pm 0.04</math></b>
PDI-Cifar100	$0.411 \pm 0.152$	$0.434 \pm 0.172$	<b><math>0.259 \pm 0.0</math></b>	$0.271 \pm 0.04$	$0.279 \pm 0.032$	<b><math>0.237 \pm 0.016</math></b>
PDI-ImageNet	$0.366 \pm 0.049$	$0.333 \pm 0.056$	<b><math>0.224 \pm 0.0</math></b>	$0.265 \pm 0.019$	$0.254 \pm 0.011$	<b><math>0.217 \pm 0.008</math></b>
PDI-LM1B	$0.665 \pm 0.016$	$0.665 \pm 0.013$	<b><math>0.65 \pm 0.002</math></b>	$0.649 \pm 0.008$	$0.655 \pm 0.013$	<b><math>0.638 \pm 0.013</math></b>
PDI-WMT	$0.446 \pm 0.09$	$0.45 \pm 0.065$	<b><math>0.372 \pm 0.0</math></b>	$0.396 \pm 0.026$	$0.398 \pm 0.019$	<b><math>0.356 \pm 0.014</math></b>
Bad Prior						
	ASHA	AsyncHB	PriorBand	ASHA	AsyncHB	PriorBand
JAHS-C10	$11.382 \pm 2.523$	$12.135 \pm 1.62$	<b><math>11.055 \pm 1.768</math></b>	<b><math>9.739 \pm 0.716</math></b>	$9.869 \pm 0.586$	$9.806 \pm 0.632$
JAHS-CH	<b><math>6.325 \pm 1.116</math></b>	$6.771 \pm 1.296$	$6.622 \pm 0.877$	$5.492 \pm 0.362$	$5.565 \pm 0.342$	<b><math>5.316 \pm 0.667</math></b>
JAHS-FM	$6.627 \pm 3.384$	$6.184 \pm 1.009$	<b><math>5.715 \pm 0.485</math></b>	<b><math>5.127 \pm 0.235</math></b>	$5.175 \pm 0.274$	$5.333 \pm 0.332$
PDI-Cifar100	<b><math>0.411 \pm 0.152</math></b>	$0.434 \pm 0.172$	$0.657 \pm 0.266$	$0.271 \pm 0.04$	$0.279 \pm 0.032$	<b><math>0.256 \pm 0.036</math></b>
PDI-ImageNet	$0.366 \pm 0.049$	$0.335 \pm 0.058$	<b><math>0.324 \pm 0.062</math></b>	$0.265 \pm 0.019$	$0.258 \pm 0.015$	<b><math>0.257 \pm 0.027</math></b>
PDI-LM1B	$0.665 \pm 0.016$	<b><math>0.665 \pm 0.013</math></b>	$0.668 \pm 0.016$	$0.649 \pm 0.008$	$0.654 \pm 0.013$	<b><math>0.646 \pm 0.012</math></b>
PDI-WMT	<b><math>0.446 \pm 0.09</math></b>	$0.45 \pm 0.065$	$0.485 \pm 0.141$	$0.396 \pm 0.026$	$0.398 \pm 0.019$	<b><math>0.375 \pm 0.014</math></b>

Table 12: Table comparing Asynchronous Successive Halving (ASHA), Asynchronous HyperBand (AsyncHB), ASHA+ $\mathcal{E}_\pi$  and Asynchronous-HyperBand+ $\mathcal{E}_\pi$  final validation errors of the current incumbent at the highest fidelity at 2 budget horizons of  $1\times$  and  $5\times$ . Runs are averaged over 10 seeds where each run is with 4 workers. This table showcases the flexibility of the ESP, as it can be applied off-the-shelf to other multi-fidelity algorithms too. Under informative priors, ASHA(+ESP) performs marginally better than Asynchronous-HB(+ESP) but wanes for longer budgets. Interestingly, ASHA(+ESP) retains a better performance under the bad priors. Since ASHA(+ESP) samples only at the  $r=0$ , the ESP fixes  $p_U$  to 50%. This could explain the increased exploration under bad priors and reduced exploitation under the good priors for ASHA(+ESP).

Benchmark	$1\times$				$5\times$			
	Near Optimum Prior							
	ASHA	AsyncHB	ASHA(+ESP)	Async-HB(+ESP)	ASHA	AsyncHB	ASHA(+ESP)	Async-HB(+ESP)
JAHS-C10	$11.382 \pm 2.523$	$12.192 \pm 1.558$	<b><math>8.281 \pm 0.221</math></b>	$8.281 \pm 0.221$	$9.739 \pm 0.716$	$9.901 \pm 0.598$	$8.223 \pm 0.183$	<b><math>8.18 \pm 0.177</math></b>
JAHS-CH	$6.325 \pm 1.116$	$6.771 \pm 1.296$	<b><math>4.466 \pm 0.188</math></b>	$4.466 \pm 0.188$	$5.492 \pm 0.362$	$5.565 \pm 0.342$	<b><math>4.466 \pm 0.188</math></b>	$4.466 \pm 0.188$
JAHS-FM	$6.627 \pm 3.384$	$6.245 \pm 0.999$	<b><math>4.713 \pm 0.052</math></b>	$4.713 \pm 0.052$	$5.127 \pm 0.235$	$5.134 \pm 0.274$	$4.709 \pm 0.049$	<b><math>4.707 \pm 0.05</math></b>
PDI-Cifar100	$0.411 \pm 0.152$	$0.434 \pm 0.172$	<b><math>0.243 \pm 0.056</math></b>	$0.259 \pm 0.083$	$0.271 \pm 0.04$	$0.279 \pm 0.032$	<b><math>0.221 \pm 0.037</math></b>	$0.227 \pm 0.04$
PDI-ImageNet	$0.366 \pm 0.049$	$0.333 \pm 0.056$	$0.269 \pm 0.07$	<b><math>0.261 \pm 0.06</math></b>	$0.265 \pm 0.019$	$0.255 \pm 0.011$	$0.22 \pm 0.021$	<b><math>0.214 \pm 0.017</math></b>
PDI-LM1B	$0.665 \pm 0.016$	$0.663 \pm 0.015$	<b><math>0.644 \pm 0.025</math></b>	<b><math>0.639 \pm 0.018</math></b>	$0.649 \pm 0.008$	$0.655 \pm 0.016$	$0.629 \pm 0.01$	<b><math>0.628 \pm 0.009</math></b>
PDI-WMT	$0.446 \pm 0.09$	$0.443 \pm 0.06$	<b><math>0.364 \pm 0.036</math></b>	$0.373 \pm 0.053$	$0.396 \pm 0.026$	$0.402 \pm 0.027$	<b><math>0.344 \pm 0.035</math></b>	$0.353 \pm 0.026$
Good Prior								
	ASHA	AsyncHB	ASHA(+ESP)	Async-HB(+ESP)	ASHA	AsyncHB	ASHA(+ESP)	Async-HB(+ESP)
JAHS-C10	$11.382 \pm 2.523$	$12.192 \pm 1.558$	<b><math>10.194 \pm 0.0</math></b>	$10.194 \pm 0.0$	$9.739 \pm 0.716$	$9.901 \pm 0.598$	$9.603 \pm 0.412$	<b><math>9.365 \pm 0.51</math></b>
JAHS-CH	<b><math>6.325 \pm 1.116</math></b>	$6.771 \pm 1.296$	$6.603 \pm 0.0$	$6.603 \pm 0.0$	$5.492 \pm 0.362$	$5.565 \pm 0.342$	$5.592 \pm 0.31$	<b><math>5.386 \pm 0.327</math></b>
JAHS-FM	$6.627 \pm 3.384$	$6.184 \pm 1.009$	<b><math>5.042 \pm 0.0</math></b>	$5.042 \pm 0.0$	$5.127 \pm 0.235$	$5.126 \pm 0.211$	<b><math>4.995 \pm 0.088</math></b>	$5.013 \pm 0.06$
PDI-Cifar100	$0.411 \pm 0.152$	$0.434 \pm 0.172$	<b><math>0.259 \pm 0.0</math></b>	$0.259 \pm 0.0$	$0.271 \pm 0.04$	$0.279 \pm 0.032$	$0.241 \pm 0.014$	<b><math>0.24 \pm 0.017</math></b>
PDI-ImageNet	$0.366 \pm 0.049$	$0.333 \pm 0.056$	<b><math>0.224 \pm 0.0</math></b>	$0.224 \pm 0.0$	$0.265 \pm 0.019$	$0.254 \pm 0.011$	$0.22 \pm 0.005$	<b><math>0.217 \pm 0.005</math></b>
PDI-LM1B	$0.665 \pm 0.016$	$0.665 \pm 0.013$	<b><math>0.649 \pm 0.004</math></b>	$0.651 \pm 0.0$	$0.649 \pm 0.008$	$0.655 \pm 0.013$	<b><math>0.646 \pm 0.006</math></b>	$0.646 \pm 0.005$
PDI-WMT	$0.446 \pm 0.09$	$0.45 \pm 0.065$	<b><math>0.372 \pm 0.0</math></b>	$0.372 \pm 0.0$	$0.396 \pm 0.026$	$0.398 \pm 0.019$	$0.371 \pm 0.004$	<b><math>0.368 \pm 0.009</math></b>
Bad Prior								
	ASHA	AsyncHB	ASHA(+ESP)	Async-HB(+ESP)	ASHA	AsyncHB	ASHA(+ESP)	Async-HB(+ESP)
JAHS-C10	<b><math>11.382 \pm 2.523</math></b>	$12.135 \pm 1.62$	$12.697 \pm 2.049$	$13.857 \pm 3.841$	$9.739 \pm 0.716$	$9.869 \pm 0.586$	<b><math>9.669 \pm 0.502</math></b>	$9.809 \pm 0.774$
JAHS-CH	<b><math>6.325 \pm 1.116</math></b>	$6.771 \pm 1.296$	$6.778 \pm 0.904$	$9.416 \pm 2.559$	<b><math>5.492 \pm 0.362</math></b>	$5.565 \pm 0.342$	$5.856 \pm 0.401$	$5.506 \pm 0.703$
JAHS-FM	$6.627 \pm 3.384$	<b><math>6.184 \pm 1.009</math></b>	$7.091 \pm 3.287$	$7.244 \pm 3.419$	$5.127 \pm 0.235$	$5.175 \pm 0.274$	<b><math>5.113 \pm 0.299</math></b>	$5.134 \pm 0.264$
PDI-Cifar100	<b><math>0.411 \pm 0.152</math></b>	$0.434 \pm 0.172$	$0.479 \pm 0.226$	$0.665 \pm 0.186$	<b><math>0.271 \pm 0.04</math></b>	$0.279 \pm 0.032$	$0.29 \pm 0.041$	$0.293 \pm 0.02$
PDI-ImageNet	$0.366 \pm 0.049$	<b><math>0.335 \pm 0.058</math></b>	$0.401 \pm 0.177$	$0.558 \pm 0.155$	$0.265 \pm 0.019$	<b><math>0.258 \pm 0.015</math></b>	$0.291 \pm 0.098$	$0.288 \pm 0.052$
PDI-LM1B	$0.665 \pm 0.016$	<b><math>0.665 \pm 0.013</math></b>	$0.687 \pm 0.046$	$0.705 \pm 0.042$	$0.649 \pm 0.008$	$0.654 \pm 0.013$	<b><math>0.647 \pm 0.018</math></b>	$0.652 \pm 0.012$
PDI-WMT	<b><math>0.446 \pm 0.09</math></b>	$0.45 \pm 0.065$	$0.54 \pm 0.172$	$0.621 \pm 0.152$	$0.396 \pm 0.026$	$0.398 \pm 0.019$	<b><math>0.383 \pm 0.022</math></b>	$0.395 \pm 0.029$

Table 13: Table comparing Bayesian Optimization (BO),  $\pi$ BO, BOHB and PriorBand+BO final validation errors of the current incumbent at the highest fidelity at 2 budget horizons of  $10\times$  and  $20\times$ . Runs are averaged over 50 seeds on 1 worker each. BO,  $\pi$ BO and PriorBand+BO use an initial design size of 10. This table highlights the model extensibility of PriorBand. Under near-optimum priors,  $\pi$ BO emerges as the best, especially after  $10\times$ . Knowledge of a near-optimum region is not available in practice and  $\pi$ BO’s rate of recovery from bad priors is costlier for DL. PriorBand+BO importantly is better than BOHB in almost all settings and better than vanilla-BO.

Benchmark	10x				20x			
	BO	$\pi$ BO	BOHB	PriorBand+BO	BO	$\pi$ BO	BOHB	PriorBand+BO
Near Optimum Prior								
JAHS-C10	9.723 $\pm$ 0.949	<b>8.23 <math>\pm</math> 0.229</b>	10.281 $\pm$ 0.88	8.252 $\pm$ 0.215	9.115 $\pm$ 0.725	<b>8.07 <math>\pm</math> 0.125</b>	9.969 $\pm$ 0.688	8.248 $\pm$ 0.21
JAHS-CH	5.445 $\pm$ 0.493	4.481 $\pm$ 0.112	5.628 $\pm$ 0.613	<b>4.452 <math>\pm</math> 0.127</b>	4.964 $\pm$ 0.426	<b>4.402 <math>\pm</math> 0.089</b>	5.393 $\pm$ 0.442	4.452 $\pm$ 0.127
JAHS-FM	5.097 $\pm$ 0.223	<b>4.71 <math>\pm</math> 0.048</b>	5.406 $\pm$ 0.323	4.724 $\pm$ 0.06	4.917 $\pm$ 0.172	<b>4.687 <math>\pm</math> 0.042</b>	5.227 $\pm$ 0.255	4.724 $\pm$ 0.06
LC-126026	0.042 $\pm$ 0.007	0.026 $\pm$ 0.006	0.043 $\pm$ 0.009	<b>0.024 <math>\pm</math> 0.007</b>	0.031 $\pm$ 0.008	<b>0.017 <math>\pm</math> 0.006</b>	0.038 $\pm$ 0.008	0.023 $\pm$ 0.005
LC-167190	0.188 $\pm$ 0.021	0.138 $\pm$ 0.014	0.191 $\pm$ 0.022	<b>0.136 <math>\pm</math> 0.016</b>	0.164 $\pm$ 0.023	<b>0.119 <math>\pm</math> 0.012</b>	0.175 $\pm$ 0.023	0.134 $\pm$ 0.014
LC-168330	0.38 $\pm$ 0.043	<b>0.289 <math>\pm</math> 0.03</b>	0.412 $\pm$ 0.032	0.295 $\pm$ 0.044	0.322 $\pm$ 0.044	<b>0.258 <math>\pm</math> 0.024</b>	0.395 $\pm$ 0.033	0.281 $\pm$ 0.035
LC-168910	0.294 $\pm$ 0.033	<b>0.198 <math>\pm</math> 0.021</b>	0.31 $\pm$ 0.033	0.2 $\pm$ 0.021	0.266 $\pm$ 0.042	<b>0.171 <math>\pm</math> 0.022</b>	0.297 $\pm$ 0.025	0.2 $\pm$ 0.02
LC-189906	0.14 $\pm$ 0.021	0.112 $\pm$ 0.011	0.141 $\pm$ 0.014	<b>0.111 <math>\pm</math> 0.012</b>	0.117 $\pm$ 0.015	<b>0.101 <math>\pm</math> 0.008</b>	0.13 $\pm$ 0.011	0.106 $\pm$ 0.011
PDI-Cifar100	0.284 $\pm$ 0.041	0.239 $\pm$ 0.035	0.291 $\pm$ 0.039	<b>0.238 <math>\pm</math> 0.035</b>	0.238 $\pm$ 0.022	<b>0.224 <math>\pm</math> 0.03</b>	0.262 $\pm$ 0.03	0.227 $\pm$ 0.029
PDI-ImageNet	0.261 $\pm$ 0.033	<b>0.222 <math>\pm</math> 0.02</b>	0.289 $\pm$ 0.035	0.238 $\pm$ 0.035	0.234 $\pm$ 0.02	<b>0.215 <math>\pm</math> 0.017</b>	0.278 $\pm$ 0.029	0.228 $\pm$ 0.028
PDI-LM1B	0.66 $\pm$ 0.015	0.644 $\pm$ 0.02	0.643 $\pm$ 0.014	<b>0.637 <math>\pm</math> 0.014</b>	0.646 $\pm$ 0.014	<b>0.631 <math>\pm</math> 0.018</b>	0.634 $\pm$ 0.014	0.633 $\pm$ 0.012
PDI-WMT	0.387 $\pm$ 0.026	<b>0.349 <math>\pm</math> 0.025</b>	0.382 $\pm$ 0.022	0.353 $\pm$ 0.027	0.362 $\pm$ 0.021	<b>0.335 <math>\pm</math> 0.022</b>	0.363 $\pm$ 0.017	0.346 $\pm$ 0.021
Good Prior								
JAHS-C10	10.049 $\pm$ 0.983	<b>9.52 <math>\pm</math> 0.427</b>	10.281 $\pm$ 0.88	9.887 $\pm$ 0.39	9.451 $\pm$ 0.806	<b>9.222 <math>\pm</math> 0.287</b>	9.969 $\pm$ 0.688	9.706 $\pm$ 0.416
JAHS-CH	5.743 $\pm$ 0.482	5.632 $\pm$ 0.323	5.628 $\pm$ 0.613	<b>5.56 <math>\pm</math> 0.51</b>	<b>5.158 <math>\pm</math> 0.387</b>	5.439 $\pm$ 0.256	5.393 $\pm$ 0.442	5.393 $\pm$ 0.417
JAHS-FM	5.257 $\pm$ 0.245	5.038 $\pm$ 0.015	5.406 $\pm$ 0.323	<b>4.906 <math>\pm</math> 0.091</b>	5.047 $\pm$ 0.205	4.965 $\pm$ 0.063	5.227 $\pm$ 0.255	<b>4.964 <math>\pm</math> 0.114</b>
LC-126026	0.049 $\pm$ 0.007	0.044 $\pm$ 0.003	0.043 $\pm$ 0.009	<b>0.039 <math>\pm</math> 0.006</b>	0.04 $\pm$ 0.008	0.04 $\pm$ 0.003	0.038 $\pm$ 0.008	<b>0.036 <math>\pm</math> 0.006</b>
LC-167190	0.204 $\pm$ 0.019	0.209 $\pm$ 0.017	0.191 $\pm$ 0.022	<b>0.184 <math>\pm</math> 0.027</b>	0.175 $\pm$ 0.023	0.203 $\pm$ 0.018	0.175 $\pm$ 0.023	<b>0.169 <math>\pm</math> 0.026</b>
LC-168330	0.412 $\pm$ 0.032	0.404 $\pm$ 0.029	0.412 $\pm$ 0.032	<b>0.374 <math>\pm</math> 0.051</b>	0.361 $\pm$ 0.04	0.366 $\pm$ 0.039	0.395 $\pm$ 0.033	<b>0.344 <math>\pm</math> 0.052</b>
LC-168910	0.318 $\pm$ 0.028	0.313 $\pm$ 0.024	0.31 $\pm$ 0.033	<b>0.302 <math>\pm</math> 0.032</b>	0.295 $\pm$ 0.03	0.301 $\pm$ 0.022	0.297 $\pm$ 0.025	<b>0.289 <math>\pm</math> 0.035</b>
LC-189906	0.153 $\pm$ 0.02	0.148 $\pm$ 0.011	0.141 $\pm$ 0.014	<b>0.126 <math>\pm</math> 0.016</b>	0.124 $\pm$ 0.013	0.141 $\pm$ 0.009	0.13 $\pm$ 0.011	<b>0.115 <math>\pm</math> 0.014</b>
PDI-Cifar100	0.282 $\pm$ 0.039	0.25 $\pm$ 0.014	0.291 $\pm$ 0.039	<b>0.248 <math>\pm</math> 0.015</b>	0.242 $\pm$ 0.028	<b>0.23 <math>\pm</math> 0.017</b>	0.262 $\pm$ 0.03	0.234 $\pm$ 0.016
PDI-ImageNet	0.266 $\pm$ 0.029	<b>0.223 <math>\pm</math> 0.005</b>	0.289 $\pm$ 0.035	0.223 $\pm$ 0.003	0.237 $\pm$ 0.017	<b>0.213 <math>\pm</math> 0.009</b>	0.278 $\pm$ 0.029	0.22 $\pm$ 0.008
PDI-LM1B	0.663 $\pm$ 0.016	0.65 $\pm$ 0.003	0.643 $\pm$ 0.014	<b>0.641 <math>\pm</math> 0.011</b>	0.648 $\pm$ 0.013	0.646 $\pm$ 0.006	<b>0.634 <math>\pm</math> 0.014</b>	0.636 $\pm$ 0.011
PDI-WMT	0.391 $\pm$ 0.018	0.372 $\pm$ 0.003	0.382 $\pm$ 0.022	<b>0.367 <math>\pm</math> 0.014</b>	0.372 $\pm$ 0.014	0.368 $\pm$ 0.009	0.363 $\pm$ 0.017	<b>0.356 <math>\pm</math> 0.018</b>
Bad Prior								
JAHS-C10	<b>10.132 <math>\pm</math> 1.143</b>	10.24 $\pm$ 1.242	10.281 $\pm$ 0.88	10.213 $\pm$ 0.846	<b>9.354 <math>\pm</math> 0.722</b>	9.534 $\pm$ 0.857	9.969 $\pm$ 0.688	9.716 $\pm$ 0.507
JAHS-CH	5.874 $\pm$ 0.803	6.572 $\pm$ 1.379	5.628 $\pm$ 0.613	<b>5.608 <math>\pm</math> 0.582</b>	<b>5.316 <math>\pm</math> 0.537</b>	6.204 $\pm$ 1.057	5.393 $\pm$ 0.442	5.429 $\pm$ 0.531
JAHS-FM	5.323 $\pm$ 0.38	5.425 $\pm$ 0.343	5.406 $\pm$ 0.323	<b>5.261 <math>\pm</math> 0.291</b>	<b>5.088 <math>\pm</math> 0.227</b>	5.212 $\pm$ 0.276	5.227 $\pm$ 0.255	5.138 $\pm$ 0.199
LC-126026	0.051 $\pm$ 0.007	0.059 $\pm$ 0.009	0.043 $\pm$ 0.009	<b>0.039 <math>\pm</math> 0.01</b>	0.042 $\pm$ 0.008	0.051 $\pm$ 0.005	0.038 $\pm$ 0.008	<b>0.034 <math>\pm</math> 0.01</b>
LC-167190	0.21 $\pm$ 0.024	0.228 $\pm$ 0.024	0.191 $\pm$ 0.022	<b>0.185 <math>\pm</math> 0.028</b>	0.184 $\pm$ 0.026	0.216 $\pm$ 0.019	0.175 $\pm$ 0.023	<b>0.163 <math>\pm</math> 0.024</b>
LC-168330	0.427 $\pm$ 0.043	0.463 $\pm$ 0.042	0.412 $\pm$ 0.032	<b>0.395 <math>\pm</math> 0.04</b>	<b>0.35 <math>\pm</math> 0.052</b>	0.429 $\pm$ 0.04	0.395 $\pm$ 0.033	0.366 $\pm$ 0.053
LC-168910	0.323 $\pm$ 0.028	0.321 $\pm$ 0.045	0.31 $\pm$ 0.033	<b>0.304 <math>\pm</math> 0.032</b>	0.291 $\pm$ 0.028	0.29 $\pm$ 0.012	0.297 $\pm$ 0.025	<b>0.288 <math>\pm</math> 0.029</b>
LC-189906	0.217 $\pm$ 0.076	0.365 $\pm$ 0.139	0.141 $\pm$ 0.014	<b>0.126 <math>\pm</math> 0.019</b>	0.129 $\pm$ 0.016	0.185 $\pm$ 0.057	0.13 $\pm$ 0.011	<b>0.115 <math>\pm</math> 0.017</b>
PDI-Cifar100	0.408 $\pm$ 0.155	0.586 $\pm$ 0.195	0.291 $\pm$ 0.039	<b>0.282 <math>\pm</math> 0.054</b>	0.251 $\pm$ 0.031	0.336 $\pm$ 0.164	0.262 $\pm$ 0.03	<b>0.241 <math>\pm</math> 0.021</b>
PDI-ImageNet	0.355 $\pm$ 0.135	0.619 $\pm$ 0.188	0.289 $\pm$ 0.035	<b>0.285 <math>\pm</math> 0.032</b>	0.26 $\pm$ 0.109	0.41 $\pm$ 0.233	0.278 $\pm$ 0.029	<b>0.249 <math>\pm</math> 0.025</b>
PDI-LM1B	0.688 $\pm$ 0.034	0.782 $\pm$ 0.06	<b>0.643 <math>\pm</math> 0.014</b>	0.644 $\pm$ 0.013	0.653 $\pm$ 0.015	0.667 $\pm$ 0.03	<b>0.634 <math>\pm</math> 0.014</b>	0.636 $\pm$ 0.012
PDI-WMT	0.523 $\pm$ 0.124	0.751 $\pm$ 0.099	0.382 $\pm$ 0.022	<b>0.38 <math>\pm</math> 0.025</b>	0.375 $\pm$ 0.024	0.497 $\pm$ 0.133	0.363 $\pm$ 0.017	<b>0.361 <math>\pm</math> 0.018</b>

Table 14: Table comparing Asynchronous-HyperBand+BO (Mobster), Mobster+ESP and PriorBand+BO final validation errors of the current incumbent at the highest fidelity at 2 budget horizons of  $1\times$  and  $5\times$ . Runs are averaged over 10 seeds where each run is with 4 workers. This table shows the extensibility of ESP to asynchronous model-based HB and also verifies the effectiveness of running PriorBand in a parallel setting. The key differences between Mobster+ESP and PriorBand+BO are their choice of initial design and the nature of multi-fidelity scheduling.

Benchmark	$1\times$			$5\times$		
<b>Good Prior</b>						
	Mobster	Mobster(+ESP)	PriorBand+BO	Mobster	Mobster(+ESP)	PriorBand+BO
JAHS-C10	$12.192 \pm 1.558$	<b><math>8.281 \pm 0.221</math></b>	$8.281 \pm 0.221$	$9.872 \pm 0.579$	<b><math>8.18 \pm 0.177</math></b>	$8.281 \pm 0.221$
JAHS-CH	$6.771 \pm 1.296$	<b><math>4.466 \pm 0.188</math></b>	$4.466 \pm 0.188$	$5.565 \pm 0.342$	$4.466 \pm 0.188$	<b><math>4.424 \pm 0.169</math></b>
JAHS-FM	$6.172 \pm 1.015$	<b><math>4.713 \pm 0.052</math></b>	$4.713 \pm 0.052$	$5.177 \pm 0.279$	<b><math>4.713 \pm 0.052</math></b>	$4.713 \pm 0.052$
PDI-Cifar100	$0.434 \pm 0.172$	$0.259 \pm 0.083$	<b><math>0.243 \pm 0.056</math></b>	$0.252 \pm 0.029$	<b><math>0.217 \pm 0.033</math></b>	$0.217 \pm 0.033$
PDI-ImageNet	$0.338 \pm 0.056$	$0.261 \pm 0.06$	<b><math>0.258 \pm 0.053</math></b>	$0.249 \pm 0.023$	<b><math>0.214 \pm 0.017</math></b>	$0.234 \pm 0.029$
PDI-LM1B	$0.663 \pm 0.012$	<b><math>0.639 \pm 0.018</math></b>	$0.642 \pm 0.022$	$0.647 \pm 0.01$	$0.634 \pm 0.012$	<b><math>0.63 \pm 0.018</math></b>
PDI-WMT	$0.45 \pm 0.065$	$0.368 \pm 0.04$	<b><math>0.358 \pm 0.032</math></b>	$0.374 \pm 0.015$	<b><math>0.348 \pm 0.023</math></b>	$0.353 \pm 0.027$
<b>Good Prior</b>						
	Mobster	Mobster(+ESP)	PriorBand+BO	Mobster	Mobster(+ESP)	PriorBand+BO
JAHS-C10	$12.192 \pm 1.558$	<b><math>10.194 \pm 0.0</math></b>	$10.194 \pm 0.0$	$9.952 \pm 0.637$	<b><math>9.297 \pm 0.39</math></b>	$9.609 \pm 0.576$
JAHS-CH	$6.768 \pm 1.295$	<b><math>6.603 \pm 0.0</math></b>	$6.603 \pm 0.0$	$5.668 \pm 0.371$	$5.492 \pm 0.312$	<b><math>5.321 \pm 0.553</math></b>
JAHS-FM	$6.245 \pm 0.999$	<b><math>5.042 \pm 0.0</math></b>	$5.042 \pm 0.0$	$5.134 \pm 0.274$	$5.013 \pm 0.06$	<b><math>5.008 \pm 0.063</math></b>
PDI-Cifar100	$0.434 \pm 0.172$	<b><math>0.259 \pm 0.0</math></b>	$0.259 \pm 0.0$	$0.252 \pm 0.029$	<b><math>0.226 \pm 0.017</math></b>	$0.234 \pm 0.02$
PDI-ImageNet	$0.333 \pm 0.056$	<b><math>0.224 \pm 0.0</math></b>	$0.224 \pm 0.0$	$0.244 \pm 0.013$	<b><math>0.222 \pm 0.005</math></b>	$0.224 \pm 0.0$
PDI-LM1B	$0.663 \pm 0.012$	$0.651 \pm 0.0$	<b><math>0.641 \pm 0.009</math></b>	$0.649 \pm 0.009$	$0.646 \pm 0.008$	<b><math>0.635 \pm 0.011</math></b>
PDI-WMT	$0.45 \pm 0.065$	<b><math>0.372 \pm 0.0</math></b>	$0.372 \pm 0.0$	$0.377 \pm 0.017$	$0.37 \pm 0.006$	<b><math>0.358 \pm 0.02</math></b>
<b>Bad Prior</b>						
	Mobster	Mobster(+ESP)	PriorBand+BO	Mobster	Mobster(+ESP)	PriorBand+BO
JAHS-C10	$12.375 \pm 1.474$	$13.857 \pm 3.841$	<b><math>10.252 \pm 1.069</math></b>	$9.792 \pm 0.624$	$9.744 \pm 0.643$	<b><math>9.614 \pm 0.622</math></b>
JAHS-CH	<b><math>6.768 \pm 1.295</math></b>	$9.218 \pm 2.769$	$7.291 \pm 1.25$	$5.668 \pm 0.371$	$5.56 \pm 0.707$	<b><math>5.407 \pm 0.483</math></b>
JAHS-FM	<b><math>6.245 \pm 0.999</math></b>	$7.244 \pm 3.419$	$6.845 \pm 3.545$	<b><math>5.134 \pm 0.274</math></b>	$5.181 \pm 0.266$	$5.17 \pm 0.227$
PDI-Cifar100	<b><math>0.434 \pm 0.172</math></b>	$0.678 \pm 0.195$	$0.657 \pm 0.266$	$0.252 \pm 0.029$	$0.258 \pm 0.025$	<b><math>0.244 \pm 0.03</math></b>
PDI-ImageNet	$0.333 \pm 0.056$	$0.558 \pm 0.155$	<b><math>0.328 \pm 0.035</math></b>	<b><math>0.244 \pm 0.013</math></b>	$0.248 \pm 0.028$	$0.266 \pm 0.026$
PDI-LM1B	$0.663 \pm 0.012$	$0.699 \pm 0.034$	<b><math>0.661 \pm 0.022</math></b>	$0.649 \pm 0.009$	$0.657 \pm 0.006$	<b><math>0.639 \pm 0.01</math></b>
PDI-WMT	<b><math>0.45 \pm 0.065</math></b>	$0.609 \pm 0.16$	$0.47 \pm 0.088$	$0.375 \pm 0.013$	$0.372 \pm 0.017$	<b><math>0.371 \pm 0.016</math></b>