

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TASK-AGNOSTIC AMORTIZED MULTI-OBJECTIVE OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Balancing competing objectives is omnipresent across disciplines, from drug design to autonomous systems. Multi-objective Bayesian optimization is a promising solution for such expensive, black-box problems: it fits probabilistic surrogates and selects new designs via an acquisition function that balances exploration and exploitation. In practice, it requires tailored choices of surrogate and acquisition that rarely transfer to the next problem, is myopic when multi-step planning is often required, and adds refitting overhead, particularly in parallel or time-sensitive loops. We present TAMO, a fully amortized, universal policy for multi-objective black-box optimization. TAMO uses a transformer architecture that operates across varying input and objective dimensions, enabling pretraining on diverse corpora and transfer to new problems without retraining: at test time, the pretrained model proposes the next design with a single forward pass. We pretrain the policy with reinforcement learning to maximize cumulative hypervolume improvement over full trajectories, conditioning on the entire query history to approximate the Pareto frontier. Across synthetic benchmarks and real tasks, TAMO produces fast proposals, reducing proposal time by 50–1000× versus alternatives while matching or improving Pareto quality under tight evaluation budgets. These results show that transformers can perform multi-objective optimization entirely in-context, eliminating per-task surrogate fitting and acquisition engineering, and open a path to foundation-style, plug-and-play optimizers for scientific discovery workflows.

1 INTRODUCTION

Multi-objective optimization (MOO; Deb et al., 2016; Gunantara, 2018) is ubiquitous in science and engineering: practitioners routinely balance accuracy vs. cost in experimental design (Schoefer et al., 2024), latency vs. quality in adaptive streaming controllers (Peroni & Gorinsky, 2025), or efficacy vs. toxicity in drug discovery (Fromer & Coley, 2023; Lai et al., 2025). In these settings, each evaluation of the black-box objectives can be slow or costly, making sample efficiency paramount; the goal is to obtain high-quality approximations of the Pareto front with a minimal number of queries.

The standard sample-efficient paradigm for such problems is Multi-objective Bayesian optimization (MOBO; Garnett, 2023): fit probabilistic surrogates for each objective, typically using Gaussian processes (GPs; Rasmussen & Williams, 2006), then select the next query by maximizing an acquisition that balances exploration–exploitation to efficiently improve a chosen multi-objective utility, such as hypervolume, scalarizations, or preference-based criteria (Daulton et al., 2020; Belakaria et al., 2019; Daulton et al., 2023b). While effective, this recipe has three drawbacks in real-world use. First, each new problem requires training surrogates from scratch and repeatedly optimizing the acquisition, adding non-trivial GP overhead that can bottleneck decision latency in parallel or time-sensitive settings. Second, performance critically depends on modeling choices (kernel, likelihood, acquisition, initialization), especially when data are scarce, a setting MOBO is intended to handle. Third, most acquisitions are myopic, optimizing a one-step gain, which can be suboptimal when Pareto-front discovery requires multi-step planning.

Amortized optimization (Finn et al., 2017; Amos et al., 2023) addresses these issues by shifting computation offline. The idea is to pre-train on a distribution of related optimization tasks, either generated synthetically or drawn from real, previously solved datasets. At test time, proposing a new design then reduces to a single forward pass. Recent efforts have explored methods for amortized

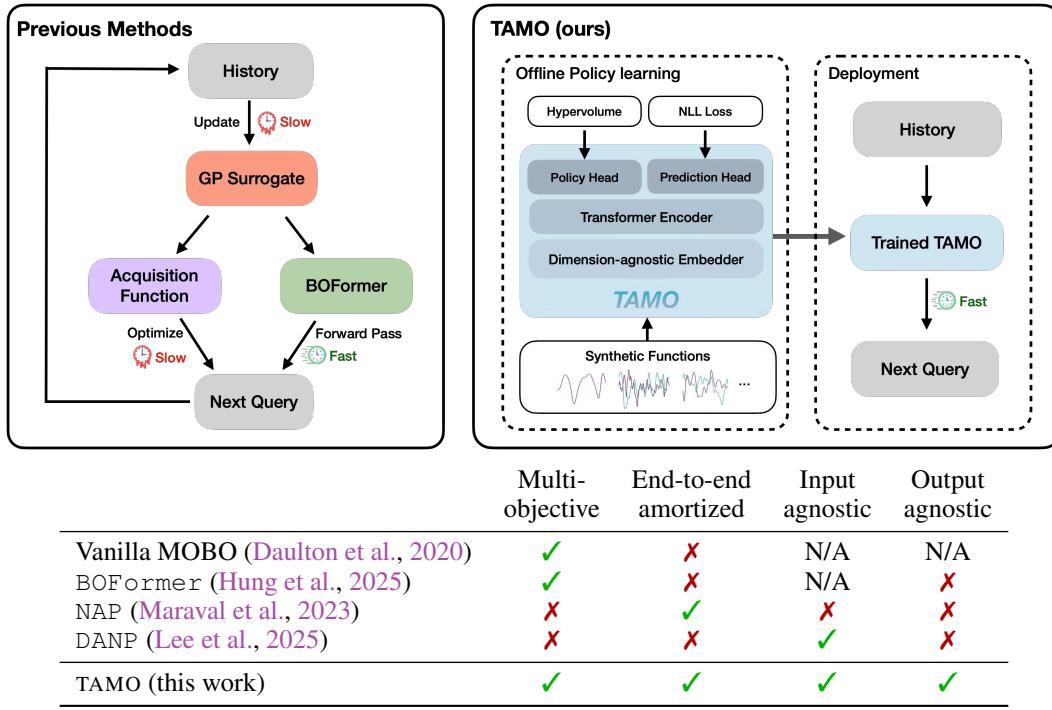


Figure 1: **Comparison of multi-objective optimization workflows.** (Top left) Previous methods like traditional MOBO or acquisition-only amortized BOFormer (Hung et al., 2025) are bottlenecked by a slow process of fitting a GP surrogate. (Top right) TAMO is *fully amortized*: a dimension-agnostic transformer policy is trained once, offline, on diverse synthetic tasks, and at deployment maps the history to the next query in a single forward pass.

Bayesian optimization (Volpp et al., 2020; Chen et al., 2022; Maraval et al., 2023; Zhang et al., 2025; Hung et al., 2025), but few address the multi-objective setting. For instance, Hung et al. (2025) only amortizes the acquisition function calculation while still relying on a GP surrogate, and its pretrained model is tied to a fixed number of objectives, which prevents transfer across heterogeneous tasks. A method that tackles these challenges would let practitioners pool heterogeneous legacy datasets for pretraining, resulting in improved outcomes in scarce-data regimes. It would also enable reusing a single optimizer as design spaces and objective counts change, and issue instant proposals in closed-loop laboratories, high-throughput campaigns, reducing overhead when evaluations are cheap or parallel.

Contributions.

- We introduce TAMO, a fully amortized policy for multi-objective optimization that maps the observed history directly to the next query (Figure 1). Training uses reinforcement learning to optimize a hypervolume-oriented utility over entire trajectories, encouraging long-horizon rather than one-step gains. At inference, proposals are produced by a single forward pass.
- TAMO is dimension agnostic on both inputs and outputs: we introduce a transformer architecture with a novel dimension-aggregating embedder that jointly represents all input features and objective values regardless of dimensionality. This enables pretraining on heterogeneous tasks, synthetic or drawn from real meta-datasets, and transfer to new problems without retraining. To our knowledge, this is the first end-to-end, dimension-agnostic architecture for black-box optimization, let alone MOO (Figure 1, bottom).
- We evaluate TAMO on synthetic and real multi-objective tasks, observing $50\times$ – $1000\times$ lower wall-clock proposal time than GP-based MOBO and baselines such as BOFORMER, which amortizes the acquisition but still relies on task-specific surrogates, while matching Pareto quality and sample efficiency. We further provide an empirical assessment of the generalization capabilities of TAMO, along with its sensitivity to deployment knobs.

108
109
110
2 PRELIMINARIES

111 **Multi-objective Optimization.** Consider a multi-objective optimization problem in which one aims
 112 to optimize a function $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), \dots, f_{d_y}(\mathbf{x})] \in \mathbb{R}^{d_y}$, and observations $\mathbf{y} = \mathbf{f}(\mathbf{x}) + \varepsilon$ where
 113 $\mathcal{X} \subset \mathbb{R}^{d_x}$ is a compact search space. In many practical scenarios, it is not possible to find a single
 114 design \mathbf{x} that is optimal for all objectives simultaneously. Instead, the notion of *Pareto dominance*
 115 is used to compare objective vectors. An objective vector $\mathbf{f}(\mathbf{x})$ *Pareto-dominates* another vector
 116 $\mathbf{f}(\mathbf{x}')$, denoted $\mathbf{f}(\mathbf{x}) \succ \mathbf{f}(\mathbf{x}')$, if $f^{(m)}(\mathbf{x}) \geq f^{(m)}(\mathbf{x}') \quad \forall m \in \{1, \dots, d_y\}$ and there exists at least
 117 one objective m' such that $f^{(m')}(\mathbf{x}) > f^{(m')}(\mathbf{x}')$. The *Pareto frontier* (PF) associated with a set of
 118 designs $X \subseteq \mathcal{X}$ is $\mathcal{P}(X) = \{\mathbf{f}(\mathbf{x}) : \mathbf{x} \in X, \nexists \mathbf{x}' \in X \text{ s.t. } \mathbf{f}(\mathbf{x}') \succ \mathbf{f}(\mathbf{x})\}$. A common goal in
 119 multi-objective optimization is to approximate the global frontier $\mathcal{P}(\mathcal{X})$ within a limited budget of T
 120 function evaluations. One popular way to assess solution quality is the *hypervolume* (HV) indicator.
 121 For a reference point $\mathbf{r} \in \mathbb{R}^{d_y}$, the hypervolume $\text{HV}(\mathcal{P}(X) \mid \mathbf{r})$ measures how much of the objective
 122 space between \mathbf{r} and the frontier $\mathcal{P}(X)$ is “covered” by Pareto-optimal points. In practice, the choice
 123 of \mathbf{r} depends on domain-specific considerations (Yang et al., 2019).

124 **Reinforcement Learning.** Recent work leverages RL to learn *non-myopic* strategies in black-box
 125 optimization, accounting for downstream impact of each evaluation (Maraval et al., 2023; Zhang et al.,
 126 2025; Hung et al., 2025). An RL problem is a Markov decision process (MDP) (Sutton et al., 1998)
 127 with states, actions, transition dynamics, a reward encoding the optimization goal, and a discount
 128 factor weighting future vs. immediate rewards. The output is a policy $\pi_\theta(\mathbf{a} \mid \mathbf{s})$, a distribution over
 129 actions \mathbf{a} for a state \mathbf{s} that maximizes expected discounted return. In the *amortized* regime, one policy
 130 is trained offline on a distribution of tasks and then deployed across new problems without retraining.

131
132 3 TASK-AGNOSTIC AMORTIZED MULTI-OBJECTIVE OPTIMIZATION (TAMO)
133

134 We introduce TAMO, a fully amortized framework for multi-objective black-box optimization. TAMO
 135 encodes the optimization history and a candidate set with a transformer backbone and directly outputs
 136 acquisition utilities. To stabilize policy learning, we additionally incorporate a prediction task into
 137 our objective function. Section 3.1 details the construction of pretraining tasks for policy learning and
 138 prediction; Section 3.2 presents the TAMO architecture; Section 3.3 formalizes the RL objective and
 139 MDP; and Section 3.4 outlines training and inference procedures. Figure 1 illustrates our workflow.

140
141 3.1 PRETRAINING DATASET CONSTRUCTION
142

143 We pre-train TAMO on a diverse distribution of synthetic multi-objective optimization tasks, denoted
 144 by $p(\tau)$. Each task $\tau \sim p(\tau)$ is defined by a black-box function $\mathbf{f}_\tau : \mathcal{X} \subset \mathbb{R}^{d_x^\tau} \rightarrow \mathbb{R}^{d_y^\tau}$, where the
 145 input and output dimensions, d_x^τ and d_y^τ , vary across tasks. This heterogeneity is key to learning a
 146 universal, dimension-agnostic policy. The full generative process for $p(\tau)$, which is based on GP
 147 priors with varied kernels and properties, is detailed in Appendix A.

148 During each training step, we sample two distinct mini-batches from task distributions to jointly
 149 optimize the model for policy learning and auxiliary prediction:

150

- 151 **Policy-learning batches.** To train the decision-making policy, each batch contains a *history set*
 152 $\mathcal{D}^h = \{(\mathbf{x}^h, \mathbf{y}^h)\}_{h=1}^{N_h}$ and a *query set* $\mathcal{D}^q = \{\mathbf{x}^q\}_{q=1}^{N_q}$. The policy conditions on the history to
 153 select the most promising query from the query set.
- 154 **Prediction batches.** To facilitate policy learning, we include an auxiliary prediction task to help
 155 the model learn the function landscape. We sample N input-output pairs from a fresh function
 156 draw. These pairs are then randomly partitioned into a *context set* $\mathcal{D}^c = \{(\mathbf{x}^c, \mathbf{y}^c)\}_{c=1}^{N_c}$ and a *target*
 157 *set* $\mathcal{D}^p = \{\mathbf{x}^p\}_{p=1}^{N_p}$, on which the model performs an in-context regression task.

158 For each dataset type $s \in \{h, q, c, p\}$, the individual elements are denoted as $x_{t,j}^s$ where $t \in$
 159 $\{1, \dots, T\}$ indexes the data point and $j \in \{1, \dots, d_x^\tau\}$ indexes the input dimension. Similarly, for
 160 outputs we have $y_{t,k}^s$ where $k \in \{1, \dots, d_y^\tau\}$ indexes the output dimension.

162 3.2 MODEL ARCHITECTURE
163

164 TAMO’s architecture is designed around a single, shared backbone that operates in two distinct tasks
165 during training: the *prediction task* and the *optimization task*. Each forward pass processes one
166 mini-batch, either a prediction batch ($\mathcal{D}^{(c)}, \mathcal{D}^{(t)}$) or an optimization batch ($\mathcal{D}^{(h)}, \mathcal{D}^{(q)}$). While the
167 input data types differ conceptually (context \leftrightarrow history and target \leftrightarrow query), they are processed by
168 the same core components. The architecture comprises four parts: (i) a *dimension-agnostic embedder*
169 mapping an observation to a vector regardless of input/output dimension; (ii) a *transformer encoder*
170 that aggregates variable-size histories/context and exposes a compact summary; (iii) lightweight
171 *task conditioning* via a small number of tokens; and (iv) two *heads*: a prediction head and a policy
172 head. The dimension-agnostic embedder and the transformer encoder blocks are shared across tasks.
173

174 **(I) Dimension-agnostic embedder.** We apply learnable scalar-to-vector maps $e_x : \mathbb{R} \rightarrow \mathbb{R}^{d_e}$ and
175 $e_y : \mathbb{R} \rightarrow \mathbb{R}^{d_e}$ dimension-wise, resulting in $e_x = e_x(\mathbf{x}) \in \mathbb{R}^{d_x^\tau \times d_e}$ and $e_y = e_y(\mathbf{y}) \in \mathbb{R}^{d_y^\tau \times d_e}$. Both
176 functions e_x and e_y are parameterized as feedforward neural networks. After L transformer
177 layers on the concatenated tokens $[e_x; e_y]$, we apply learnable dimension-specific positional tokens
178 $p_x \in \mathbb{R}^{d_x^\tau \times d_e}$ and $p_y \in \mathbb{R}^{d_y^\tau \times d_e}$ element-wise and mean-pool across the $d_x^\tau + d_y^\tau$ token axis to obtain
179 a single representation $\mathbf{E} \in \mathbb{R}^{d_e}$. These positional tokens are randomly sampled for each batch
180 from fixed pools of learned embeddings. We introduce the positional tokens to prevent the spurious
181 symmetries over dimensionalities from a permutation-invariant set encoder, allowing the model to
182 distinguish between features and objectives with the same values. During training, the embedder is
183 applied to \mathcal{D}^h and \mathcal{D}^q to yield \mathbf{E}^h and \mathbf{E}^q for the *optimization task*, and to \mathcal{D}^c and \mathcal{D}^p to yield \mathbf{E}^c
184 and \mathbf{E}^p for the *prediction task*. Each observation contributes $\mathcal{O}(1)$ tokens, so the cost scales with the
185 number of observations, not with d_x^τ or d_y^τ . Figure 2 summarizes the embedder.
186

187 **(II) Transformer encoder-decoder.** We stack $B := B_1 + B_2$ transformer layers and split them into two phases. For the *first*
188 B_1 *layers*, the *observed* tokens interact. The history (or context)
189 tokens undergo self-attention to produce $\hat{\mathbf{E}}^h$ (or $\hat{\mathbf{E}}^c$), capturing
190 intra-set structure. The query (or target) tokens then use cross-
191 attention with the keys/values provided by $\hat{\mathbf{E}}^h$ (or $\hat{\mathbf{E}}^c$), yielding
192 $\hat{\mathbf{E}}^q$ (or $\hat{\mathbf{E}}^p$). No task-specific tokens are present in B_1 . This
193 phase is the *only* path through which queries/targets access
194 information from the history/context. Then, for the *last* B_2 *layers*,
195 the sequence is reduced to *only* the query/target tokens
196 together with a small set of task-specific tokens (defined below).
197 All history/context tokens are removed from the sequence. An
198 attention mask enforces that, in these final layers, query/target
199 tokens are permitted to attend *only* to the task-specific tokens
200 (no query–query attention and no access to history/context).
201 The task-specific tokens may self-attend among themselves.
202

203 **(III) Task-specific tokens.** Task-specific tokens are *introduced*
204 *only at the entrance to the last* B_2 *layers*. For each task type, we introduce distinct tokens that guide
205 the decoder’s computation. For ***prediction tasks***, the additional tokens comprise a *prediction task*
206 token and the *output-index* positional token $p_y^{(k)}$ indicating which scalar $y_{i,k}^p$ is to be predicted. For
207 ***optimization tasks***, the additional tokens comprise an *optimization task* token, a time-budget token
208 $\mathbf{g}_{\text{time}} = \text{MLP}_\theta((T-t)/T)$, and an aggregated *input-dimension* token $\sum_{j=1}^{d_x^\tau} p_x^{(j)}$. An attention mask
209 restricts query/target tokens in B_2 to attend only to these tokens; the history/context never appears in
210 B_2 . This design yields constant overhead in d_x^τ and d_y^τ , and linear cost in token size.
211

212 **(IV) Heads.** The architecture terminates in two heads that share the backbone but produce different
213 outputs depending on the task.
214

215 ***Prediction head.*** Given the prediction tokens $\{\hat{\mathbf{E}}_i^p\}_{i=1}^{N_p}$ and an output-index positional token $p_y^{(k)}$, the
216 model produces, for each prediction input \mathbf{x}_i^p , the parameters of a K -component univariate Gaussian
217 mixture that models the scalar $y_{i,k}^p$. Concretely, an MLP applied to $\hat{\mathbf{E}}_i^p$ yields mixture weights
218 $\{\phi_{i\ell}\}_{\ell=1}^K$, means $\{\mu_{i\ell}\}_{\ell=1}^K$, and positive scales $\{\sigma_{i\ell}\}_{\ell=1}^K$, with the weights normalized by a softmax
219 ($\sum_{\ell=1}^K \phi_{i\ell} = 1$) and the scales enforced positive via a softplus transform. The resulting predictive
220

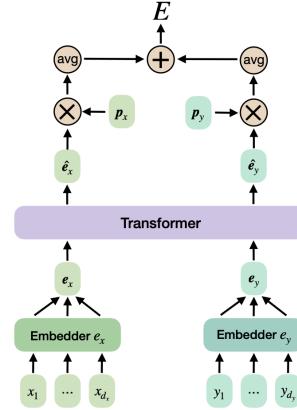


Figure 2: Dimension-agnostic embedder for a single observation.

216 density is:

$$218 \quad p(y_{i,k}^p \mid \mathbf{x}_i^p, \mathcal{D}^c) = \sum_{\ell=1}^K \phi_{i\ell} \mathcal{N}(y_{i,k}^p; \mu_{i\ell}, \sigma_{i\ell}^2). \quad (1)$$

220 Prediction tasks use samples disjoint from optimization tasks to prevent reward leakage.

221 *Policy head.* Given the query tokens $\{\hat{\mathbf{E}}_i^q\}_{i=1}^{N_q}$, the model assigns each query \mathbf{x}_i^q a scalar acquisition
222 utility $\alpha_i := \text{MLP}_\theta(\hat{\mathbf{E}}_i^q)$. These utilities are converted into a categorical policy over the query set
223 using a softmax:

$$225 \quad \pi_\theta(\mathbf{x}_i^q \mid t, T, \mathcal{H}_{1:t-1}) = \frac{\exp(\alpha_i)}{\sum_{r=1}^{N_q} \exp(\alpha_r)}. \quad (2)$$

227 During training we sample actions from this distribution; at inference we act greedily by selecting the
228 next query with the largest probability. [For a detailed visualization of the full architecture, please
229 refer to Figure S1.](#)

231 3.3 POLICY LEARNING

232 We cast the dimension-agnostic optimization problem as a Markov decision process (MDP):

- 234 • **State.** At step t , the state is $s_t = (\mathcal{D}^h, t, T)$, where \mathcal{D}^h is the current historical observations and
235 T is the total budget.
- 236 • **Action.** The action selects a candidate index $i^* \in \{1, \dots, N_q\}$ and sets $\mathbf{x}_t = \mathbf{x}_{i^*}^q$.
- 237 • **Reward.** After querying the objectives and observing \mathbf{y}_t , we update the history $\mathcal{D}^h := \mathcal{D}^h \cup$
238 $\{(\mathbf{x}_t, \mathbf{y}_t)\}$ and define the *normalized hypervolume level*:

$$239 \quad r_t = \frac{\text{HV}(\mathcal{P}(\mathcal{D}^h) \mid \mathbf{r})}{\text{HV}_\tau^*}, \quad \text{HV}_\tau^* := \text{HV}(\mathcal{P}(\mathcal{X}) \mid \mathbf{r}), \quad \text{HV}(\mathcal{P}(\emptyset) \mid \mathbf{r}) = 0.$$

242 Here HV_τ^* is the task-wise hypervolume of the optimal frontier with respect to the fixed ref-
243 erence point \mathbf{r} . We set \mathbf{r} to the componentwise worst value, $\mathbf{r} = [\hat{y}^{(1)}, \dots, \hat{y}^{(d_y)}]$ with
244 $\hat{y}^{(k)} := \min_{\mathbf{x} \in \mathcal{X}} f_\tau^{(k)}(\mathbf{x})$, which makes the hypervolume well defined and the reward bounded in
245 $[0, 1]$. This ratio measures the fraction of maximum achievable hypervolume already captured by
246 the current Pareto approximation (larger is better), and the normalization provides scale invariance
247 across heterogeneous tasks (Teoh et al., 2025).

248 The policy $\pi_\theta(\mathbf{x} \mid s)$ maximizes the expected discounted return:

$$249 \quad J(\theta) = \mathbb{E}_{\tau \sim p(\tau)} \left[\mathbb{E}_{\pi_\theta} \left[\sum_{t=1}^T \gamma^{t-1} r_t \right] \right], \quad (3)$$

252 and we estimate gradients with REINFORCE (Williams, 1992):

$$254 \quad \nabla_\theta J(\theta) \approx \mathbb{E}_{\tau, \pi_\theta} \left[\sum_{t=1}^T \nabla_\theta \log \pi_\theta(\mathbf{x}_t \mid s_t) \gamma^{t-1} r_t \right]. \quad (4)$$

256 Because we train on synthetic tasks, we can evaluate r_t exactly during learning.

258 3.4 TRAINING AND INFERENCE

260 **Training.** We train TAMO in two phases. First, we warm up the backbone on the prediction task by
261 minimizing a negative log-likelihood over $(\mathcal{D}^c, \mathcal{D}^p)$, which encourages accurate in-context regression
262 and useful representations:

$$263 \quad \mathcal{L}^{(p)}(\theta) = -\mathbb{E}_{\tau \sim p(\tau)} \left[\frac{1}{N_p d_y^p} \sum_{i=1}^{N_p} \sum_{k=1}^{d_y^p} \log p(y_{i,k}^p \mid \mathbf{x}_i^p, \mathcal{D}^c) \right]. \quad (5)$$

266 After warm-up we transition to the joint training phase, where we optimize the policy with the
267 trajectory objective $J(\theta)$ (Eq. 3), aligning the learning signal with improvements in Pareto quality
268 alongside the prediction objective. The overall objective combines both terms:

$$269 \quad \mathcal{L}(\theta) = \lambda_p \mathcal{L}^{(p)}(\theta) + \mathcal{L}^{(\text{rl})}(\theta), \quad \mathcal{L}^{(\text{rl})}(\theta) = -J(\theta), \quad (6)$$

270 and is optimized with REINFORCE (Eq. 4); the coefficient $\lambda_p > 0$ trades off prediction and policy
 271 signals. In all experiments, we fixed $\lambda_p = 1.0$. Specifically, the prediction loss $\mathcal{L}^{(p)}$ and RL loss $\mathcal{L}^{(rl)}$
 272 are calculated from two distinct forward passes through the model with different datasets, which are
 273 then summed for a single backward pass. Training on full trajectories directly rewards long-horizon
 274 improvements, while amortization enables learning from many tasks offline.

275 **Inference.** At deployment, TAMO iteratively approximates the Pareto frontier under a budget T . We
 276 initialize the history with a random observation $\mathcal{D}^h \leftarrow \{\mathbf{x}_0^h, y_0^h\}$ and set $t \leftarrow 1$. Each iteration scores
 277 the current candidate set \mathcal{D}^q with a single forward pass and proposes:
 278

$$279 \quad \mathbf{x}_t = \arg \max_{\mathbf{x}_i^q \in \mathcal{D}^q} \pi_{\theta}(\mathbf{x}_i^q \mid t, T, \mathcal{D}^h). \quad (7)$$

282 The proposed query \mathbf{x}_t is then evaluated, and the resulting observation is used to update the history.
 283 This process is iterated until the cumulative evaluation cost meets the budget. Detailed descriptions
 284 of the algorithms for training and inference are provided in Appendix B .
 285

288 4 RELATED WORK

290 **Multi-objective Bayesian Optimization (MOBO).** This line of work builds on Bayesian Optimization
 291 (Garnett, 2023), leveraging a combination of a statistical surrogate and an acquisition
 292 function to seek high-quality approximations to the Pareto set under tight evaluation budgets. Three
 293 families are prominent. *scalarization* methods convert MOBO into single-objective subproblems
 294 (e.g., ParEGO, TS-TCH; Knowles, 2006; Paria et al., 2020), letting practitioners reuse mature BO
 295 tooling and sweep preferences in parallel. *Indicator-based* methods optimize hypervolume-oriented
 296 criteria such as EHVI or HVKG (Daulton et al., 2020; 2023a), directly aligning the acquisition with
 297 the final Pareto-quality metric. Lastly, *information-theoretic* methods (PESMO, MESMO, PFES;
 298 Hernández-Lobato et al., 2016; Belakaria et al., 2021; Suzuki et al., 2020) select points that maximize
 299 information gain about the Pareto set or frontier, offering a principled exploration strategy. These
 300 approaches are effective but hinge on a carefully tuned, task-specific surrogate–acquisition pairing
 301 that must be refit and re-optimized at each iteration, all while remaining largely myopic. We instead
 302 learn a fully amortized policy that reduces design proposal to a single neural-network forward pass,
 303 dramatically lowering inference latency.

304 **Amortization and meta-learning.** Amortization replaces per-task inference with a model trained
 305 offline to operate *in-context*, exemplified by prior-data fitted transformers that achieve strong in-
 306 context performance after pretraining on large, heterogeneous datasets (Hollmann et al., 2025;
 307 Qu et al., 2025). In parallel, Conditional Neural Processes and their transformer variants learn
 308 predictors that condition on a context set and generalize via a single forward pass (Garnelo et al.,
 309 2018; Kim et al., 2019; Nguyen & Grover, 2022; Chang et al., 2025), with recent work extending
 310 them to *dimension-agnostic* settings (Dutordoir et al., 2023; Lee et al., 2025). These works focus on
 311 amortizing *prediction*. More recently, several studies leverage in-context pretrained neural processes
 312 for sequential decision-making (Huang et al., 2024; Zhang et al., 2025; Huang et al., 2025); our
 313 approach falls into this line as well.

314 **Amortized black-box optimization.** Several approaches train neural networks to amortize black-box
 315 optimization directly, typically by mapping histories to proposals or by predicting acquisition values,
 316 with success under scalar observations (Volpp et al., 2020; Chen et al., 2022; Yang et al., 2023;
 317 Maraval et al., 2023; Song et al., 2024; Huang et al., 2024), and even binary feedback (Zhang et al.,
 318 2025). Complementary to surrogate/acquisition amortization, transfer-BO with Monte Carlo Tree
 319 Search learns the search space itself by building a data-driven hierarchy of promising subregions on
 320 source tasks and reusing it to warm-start a new target before adapting online (Wang et al., 2024).
 321 However, none of these methods is simultaneously end-to-end (no per-task surrogate or acquisition),
 322 natively multi-objective, and capable of cross-dimensional transfer. The closest is BOFormer (Hung
 323 et al., 2025), which uses sequence modeling to mitigate myopia in MOBO, but still relies on task-
 324 specific surrogates and fixed output dimensional setups, necessitating additional training when the
 325 dimension changes. We address all three by pretraining a fully amortized, foundation-style policy.

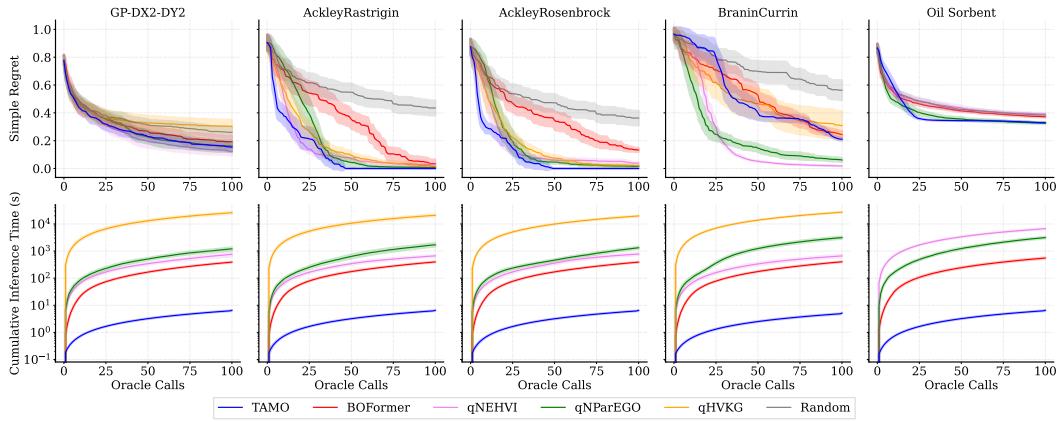


Figure 3: **Synthetic and real-world multi-objective benchmarks:** simple regret (top) and cumulative inference time (bottom) vs. oracle calls (mean \pm 95% CIs over 30 runs). **TAMO achieves competitive regret while cutting proposal time by 50 \times –1000 \times .**

5 EXPERIMENTS

We evaluate TAMO on synthetic GP tasks and standard analytic testbeds, as well as on real-world problems (Section 5.1). Subsequently, Section 5.2 studies the generalization capabilities of TAMO, e.g., with respect to unseen task dimensionalities during training on both synthetic tasks and a real-world problem. We conclude with several ablation studies related to the batch size and query set size employed at inference time (Section 5.3). Additional experiments can be found in Appendix E. **We emphasize that a single pretrained model is used across all experiments.**

Baselines. We compare against strong MOBO baselines, including decomposition and indicator-based methods (qNParEGO [Knowles 2006](#), qNEHVI [Daulton et al. 2020](#), qHVKG [Daulton et al. 2023a](#)), sequence-modeling MOBO (BOFormer [Hung et al. 2025](#)), and a random search baseline. Baselines are tuned with their recommended defaults unless otherwise noted.

Metrics. We report performance via HV-based simple regret at a fixed evaluation budget. We also measure *wall-clock* proposal time end-to-end, which for GP-based baselines includes surrogate fitting and acquisition optimization, and for our method consists of a single forward pass. For single-objective, we additionally report standard simple regret.

Implementation. TAMO is implemented using PyTorch ([Paszke et al., 2019](#)). Hyperparameter settings can be found in Appendix D.3. Code will be made available upon acceptance. For all vanilla MOBO baselines, we used the implementation from the BoTorch library ([Balandat et al., 2020](#)). For BOFormer ([Hung et al., 2025](#)), we used the publicly available implementation and pretrained model from [its official code repository](#). To ensure a fair comparison, the domain size (i.e., the size of the candidate query set) during testing is set to 2048, consistent with the configuration used for TAMO.

5.1 SYNTHETIC AND REAL-WORLD TASKS

Synthetic examples. On synthetic MOO testbeds (details in Section D.2), TAMO attains competitive or better simple regret across the entire budget (Figure 3). On GP-DX2-DY2, which is in-distribution for all methods (30 GP draws), TAMO performs on par with the best GP baselines. On the remaining three problems, *out-of-distribution* for all baselines, TAMO yields the strongest performance, except on Branin–Currin where qNEHVI and qNParEGO do better. We hypothesize this gap stems from the objectives in being well described by long length scales, outside the reach of our pretraining corpus: synthetic GP samples using lengthscales $\ell \sim \mathcal{N}(2/3, 0.5)$ over $[-5, 5]^{d_x}$ (Section D.3). Lastly, our method can also be applied effortlessly to single-objective BO, yielding competitive results compared to other GP-based alternatives (Figure S2).

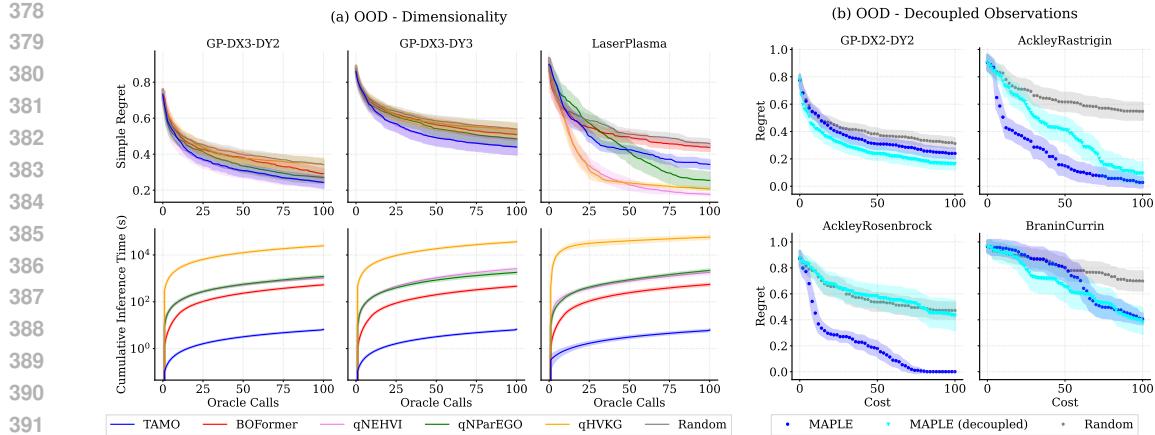


Figure 4: **Out-of-distribution evaluations.** (a) **Dimensionality:** simple regret (top) and cumulative inference time (bottom) on tasks whose input/output dimensions are unseen at pretraining. (b) **Decoupled observations:** regret vs. *cumulative cost* when, at step t , the optimizer may observe both objectives at cost 2 (dark blue) or only one at cost 1 (cyan). Curves show means with 95% confidence intervals over 60 runs (GP-DX3-DY2, GP-DX3-DY3) and 30 runs (others) with random initial observations. **TAMO shows reasonable generalization across unseen dimensionalities and decoupled feedback settings, while retaining orders-of-magnitude faster proposal times and broadly competitive regret**

Real-world example. We compare our model, pretrained only with synthetic GP samples, with other baselines on the real-world oil sorbent multi-objective problem (Daulton et al., 2022). The result is shown in Figure 3. TAMO remains competitive with GP-based alternatives, yielding the best performance, closely followed by qNParEGO.

Wall-clock time. Nevertheless, the primary advantage is speed: cumulative inference time is lower by roughly 50x–1000x, growing slowly with budget because each proposal is a single forward pass. By contrast, GP-based methods incur substantial overhead from repeated surrogate refits and acquisition optimization. Even BOFormer, which amortizes the acquisition but still relies on a GP surrogate, remains noticeably slower than TAMO.

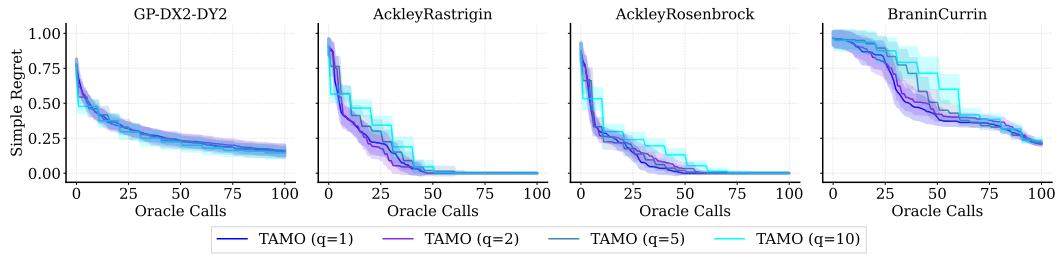
5.2 GENERALIZATION

We investigate the generalization capabilities of TAMO in two different test-time scenarios: unseen dimensionalities, or decoupled observations.

Out-of-distribution dimensionalities. We test cross-dimensional transfer by pretraining TAMO on GP tasks with $d_x \in \{1, 2\}$ and $d_y \in \{1, 2, 3\}$, then evaluating on (i) GP-DX3-DY2 and GP-DX3-DY3, and (ii) the real-world LaserPlasma task ($d_x = 4$, $d_y = 3$; Section D.2), all with unseen input/output dimensionalities. On the synthetic OOD settings (Figure 4a, left, middle), TAMO attains regret broadly comparable to the strongest GP baselines across the budget; even with 60 repetitions, we do not observe statistically decisive differences between methods. On LaserPlasma (Figure 4a, right), TAMO improves over BOFormer (which amortizes only the acquisition) but trails conventional MOBO baselines in regret. Across all cases, TAMO retains orders-of-magnitude advantages in cumulative inference time.

Decoupled observations. We next test generalization to *decoupled* settings, where objectives can be measured independently, a common setting when jointly observing all objectives is infeasible or costly, also arising when historical logs contain partial objective labels. Budget $T=100$ with cost 1 per objective: a full evaluation costs d_y , a single-objective probe costs 1. Hence, a coupled policy can do at most T/d_y full evals, while a *decoupled* one can take up to T single-objective measurements. Figure 4b plots regret vs. cumulative cost. On GP-DX2-DY2, Ackley–Rastrigin, and Branin–Currin, the *decoupled* variant of TAMO closely tracks the coupled policy, indicating that

432 TAMO can accommodate partial-feedback acquisition without retraining, offering a flexible trade-
 433 off between measurement cost and optimization progress. The exception is Ackley–Rosenbrock,
 434 where decoupling hurts performance, likely because the objectives peak at disparate locations, so
 435 single-objective measurements transfer poorly and bias the search toward one goal.
 436



437 **Figure 5: Effect of batch size on synthetic problems:** simple regret for TAMO with $q \in \{1, 2, 5, 10\}$.
 438 Curves show means with 95% CIs over 30 runs. **Smaller q converges fastest; larger q incurs a**
 439 **mild slowdown, compatible with wall-clock savings for parallel evaluations.**

440 5.3 ABLATION STUDY

441 We examine TAMO’s sensitivity to deployment knobs: *batch size* (q) and *query set size* (N_q), and
 442 quantify the accuracy–latency trade-offs they induce.

443 **Batch size.** We compare TAMO with $q \in \{1, 2, 5, 10\}$. For $q > 1$ we form batches via a light
 444 *fantasy loop* (Chang et al., 2022): pick $\mathbf{x}_t = \text{argmax}_{\mathbf{x}_i^q \in \mathcal{D}^q} \pi_\theta(\mathbf{x}_i^q \mid t, T, \mathcal{D}^h)$, predict a provisional
 445 outcome \hat{y}_t , augment $\mathcal{D}^h \leftarrow \mathcal{D}^h \cup (\mathbf{x}_t, \hat{y}_t)$, and repeat until q points are chosen (i.e., q forward
 446 passes). Across all problems, smaller batches reduce simple regret fastest; larger q slows progress
 447 modestly, consistent with the lack of real feedback within a batch (hallucination). Nonetheless,
 448 degradation remains limited and all settings improve steadily with budget, indicating that when
 449 parallel evaluations are cheap, $q > 1$ can cut wall-clock time for a small accuracy cost (Figure 5).

450 **Query set size.** TAMO scores N_q candidates per step; larger N_q improves coverage but raises proposal
 451 cost. For $N_q \in \{256, 512, 1024, 2048\}$, regret is largely insensitive, except for Branin–Currin, where
 452 a small N_q misses promising regions (Figure S3). Cumulative inference time grows roughly linearly
 453 with N_q . Even for $N_q = 2048$ (our default), TAMO remains much faster than other baselines.

466 6 DISCUSSION

467 We introduced TAMO, a fully amortized, task-agnostic policy for multi-objective black-box optimization.
 468 A single transformer backbone, trained offline with a prediction warm-up and a policy-level RL
 469 objective, maps histories to proposals in one forward pass and operates across varying input/output di-
 470 mensionalities. Empirically, TAMO delivers proposal times 50 to 1000 times lower than conventional
 471 baselines while matching Pareto quality under tight evaluation budgets, sometimes even improving it.

472 **Limitations.** Our study highlights two main axes. First, pretraining data composition: although
 473 synthetic GP corpora provide scale and control, they may miss salient real-world structure. Down-
 474 stream performance is likely sensitive to GP kernel families and smoothness, input metrics (e.g.,
 475 Mahalanobis/rotated anisotropy), multi-output correlations (coregionalization models), observation
 476 models (homo-/heteroscedastic noise, decoupled observations), and simple landscape priors (e.g.,
 477 adding a quadratic “bowl”). A systematic analysis that varies these ingredients would clarify how
 478 synthetic dataset design drives transfer. Second, inference currently assumes a discrete candidate
 479 pool, which can be restrictive in high-dimensional design spaces and in generative settings (e.g., *de*
 480 *novo* drug design) where the action space is continuous or combinatorial. Nevertheless, in pool-based
 481 scenarios like high-throughput screening and library/catalog search, this assumption aligns with
 482 practice, and our method is highly effective.

483 **Perspectives.** We envision a bright future for TAMO as a *universal engine* for black-box optimization.
 484 The modular design invites extensions to black-box constraints, cost-aware and multi-fidelity settings,

486 while retaining the single-pass interface. Beyond scalar observations, alternative data modalities, like
 487 preferential and multi-fidelity feedback, can also be incorporated with minor architecture changes. A
 488 key challenge is scaling to higher-dimensional spaces, with promising directions including factorizing
 489 the policy across input dimensions [and moving from pool-based scoring to continuous policies or](#)
 490 [generative proposal mechanisms, in the spirit of amortized design networks \(Foster et al., 2021\).](#)
 491 Lastly, as mentioned above, further work will investigate how the composition of the synthetic pre-
 492 training corpus influences downstream performance, an important direction for improving robustness
 493 and out-of-distribution behavior of amortized BO policies. Together, these advances position TAMO
 494 to serve as a foundation model-style optimizer that transfers across domains, objectives, and design
 495 spaces with minimal per-task tuning.

496 **ETHICS STATEMENT**

497 All authors have read and will adhere to the ICLR Code of Ethics. This work does not involve human
 498 subjects, personally identifiable information, or sensitive attributes; experiments use synthetic data
 499 and standard public benchmarks. We are not aware of foreseeable harms from the methodology
 500 beyond typical risks of algorithmic misuse; the intended use is scientific and engineering optimization.
 501 Compute and environmental impact were kept reasonable (single-model pretraining and standard
 502 hardware); we report settings to support reproducibility. We will respect licenses of any third-party
 503 assets used and disclose any conflicts of interest if they arise.

504 **REPRODUCIBILITY STATEMENT**

505 We document all experimental settings needed to facilitate replication: hyperparameters and optimizer
 506 details (Section D), procedures for pretraining dataset generation (Section A), and step-by-step
 507 algorithms for the training and prediction workflows (Section B). Code will be made available upon
 508 acceptance.

512 **REFERENCES**

514 Brandon Amos et al. Tutorial on Amortized Optimization. *Foundations and Trends in Machine*
 515 *Learning*, 2023.

517 Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson,
 518 and Eytan Bakshy. BoTorch: A framework for efficient monte-carlo bayesian optimization. In
 519 *Advances in Neural Information Processing Systems*, 2020.

520 Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value Entropy Search for
 521 Multi-Objective Bayesian Optimization. In *Advances in Neural Information Processing Systems*,
 522 2019.

524 Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Output space entropy search framework
 525 for multi-objective bayesian optimization. *Journal of artificial intelligence research*, 2021.

526 Paul E Chang, Prakhar Verma, ST John, Victor Picheny, Henry Moss, and Arno Solin. Fantasizing
 527 with dual gps in bayesian optimization and active learning. *arXiv preprint arXiv:2211.01053*,
 528 2022.

530 Paul Edmund Chang, Nasrulloh Ratu Bagus Satrio Loka, Daolang Huang, Ulpur Remes, Samuel Kaski,
 531 and Luigi Acerbi. Amortized probabilistic conditioning for optimization, simulation and inference.
 532 In *International Conference on Artificial Intelligence and Statistics*, pp. 703–711. PMLR, 2025.

533 Yutian Chen, Xingyou Song, Chansoo Lee, Zi Wang, Richard Zhang, David Dohan, Kazuya
 534 Kawakami, Greg Kochanski, Arnaud Doucet, Marc’Aurelio Ranzato, et al. Towards learning uni-
 535 versal hyperparameter optimizers with transformers. *Advances in Neural Information Processing*
 536 *Systems*, 35:32053–32068, 2022.

538 Sam Daulton, Maximilian Balandat, and Eytan Bakshy. Hypervolume knowledge gradient: a looka-
 539 head approach for multi-objective bayesian optimization with partial information. In *International*
Conference on Machine Learning, 2023a.

540 Sam Daulton, Maximilian Balandat, and Eytan Bakshy. Hypervolume knowledge gradient: A looka-
 541 head approach for multi-objective Bayesian optimization with partial information. In *Proceedings*
 542 *of the 40th International Conference on Machine Learning*, 2023b.

543

544 Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable Expected Hypervolume Im-
 545 provement for Parallel Multi-Objective Bayesian Optimization. In *Advances in Neural Information*
 546 *Processing Systems*, 2020.

547 Samuel Daulton, Xingchen Wan, David Eriksson, Maximilian Balandat, Michael A Osborne, and
 548 Eytan Bakshy. Bayesian optimization over discrete and mixed spaces via probabilistic repara-
 549 meterization. *Advances in Neural Information Processing Systems*, 35:12760–12774, 2022.

550

551 Kalyanmoy Deb, Karthik Sindhya, and Jussi Hakanen. Multi-objective optimization. In *Decision*
 552 *sciences*, pp. 161–200. CRC Press, 2016.

553 Vincent Dutordoir, Alan Saul, Zoubin Ghahramani, and Fergus Simpson. Neural diffusion processes.
 554 In *International Conference on Machine Learning*, pp. 8990–9012. PMLR, 2023.

555

556 Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of
 557 deep networks. In *International Conference on Machine Learning (ICML)*, pp. 1126–1135, 2017.

558 Adam Foster, Desi R Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing
 559 sequential bayesian experimental design. In *International conference on machine learning*, 2021.

560

561 Jenna C. Fromer and Connor W. Coley. Computer-aided multi-objective optimization in small
 562 molecule discovery. *Patterns*, 4(2):100678, 2023. ISSN 2666-3899. doi: <https://doi.org/10.1016/j.patter.2023.100678>. URL <https://www.sciencedirect.com/science/article/pii/S2666389923000016>.

562

563

564

565 Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray
 566 Shanahan, Yee Whye Teh, Danilo Rezende, and S. M. Ali Eslami. Conditional Neural Processes.
 567 In *International Conference on Machine Learning*, 2018.

568

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

570

571 Nyoman Gunantara. A review of multi-objective optimization: Methods and its applications. *Cogent*
 572 *Engineering*, 5(1):1502242, 2018.

573

574 Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. Predictive entropy
 575 search for multi-objective bayesian optimization. In *International conference on machine learning*,
 2016.

576

577 Noah Hollmann, Samuel Müller, Lennart Purucker, Arjun Krishnakumar, Max Körfer, Shi Bin Hoo,
 578 Robin Tibor Schirrmeister, and Frank Hutter. Accurate predictions on small data with a tabular
 579 foundation model. *Nature*, 637(8045):319–326, 2025.

580

581 Daolang Huang, Yujia Guo, Luigi Acerbi, and Samuel Kaski. Amortized bayesian experimental
 582 design for decision-making. *Advances in Neural Information Processing Systems*, 37:109460–
 109486, 2024.

583

584 Daolang Huang, Xinyi Wen, Ayush Bharti, Samuel Kaski, and Luigi Acerbi. Aline: Joint amortization
 585 for bayesian inference and active data acquisition. *arXiv preprint arXiv:2506.07259*, 2025.

586

587 Yu Heng Hung, Kai-Jie Lin, Yu-Heng Lin, Chien-Yi Wang, Cheng Sun, and Ping-Chun Hsieh.
 588 BOFormer: Learning to Solve Multi-Objective Bayesian Optimization via Non-Markovian RL. In
 589 *International Conference on Learning Representations*, 2025.

590

591 Faran Irshad, Stefan Karsch, and Andreas Döpp. Multi-objective and multi-fidelity Bayesian opti-
 592 mization of laser-plasma acceleration. *Physical Review Research*, 5(1):013063, 2023.

593

594 Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol
 595 Vinyals, and Yee Whye Teh. Attentive Neural Processes. In *International Conference on Learning*
 596 *Representations*, 2019.

594 Joshua Knowles. Parego: A hybrid algorithm with on-line landscape approximation for expensive
 595 multiobjective optimization problems. *IEEE transactions on evolutionary computation*, 2006.
 596

597 Helen Lai, Christos Kannas, Alan Kai Hassen, Emma Granqvist, Annie M. Westerlund, Djork-Arné
 598 Clevert, Mike Preuss, and Samuel Genheden. Multi-objective synthesis planning by means of
 599 monte carlo tree search. *Artificial Intelligence in the Life Sciences*, 2025.

600 Hyungi Lee, Chaeyun Jang, Dongbok Lee, and Juho Lee. Dimension agnostic neural processes.
 601 *International Conference on Learning Representations*, 2025.

602 Alexandre Maraval, Matthieu Zimmer, Antoine Grosnit, and Haitham Bou Ammar. End-to-end
 603 Meta-Bayesian Optimisation with Transformer Neural Processes. *Advances in Neural Information
 604 Processing Systems*, 36:11246–11260, 2023.

605

606 Tung Nguyen and Aditya Grover. Transformer neural processes: Uncertainty-aware meta learning
 607 via sequence modeling. In *International Conference on Machine Learning*, 2022.

608 Biswajit Paria, Kirthevasan Kandasamy, and Barnabás Póczos. A flexible framework for multi-
 609 objective bayesian optimization using random scalarizations. In *Uncertainty in Artificial Intelli-
 610 gence*, 2020.

611

612 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 613 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward
 614 Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
 615 Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep
 616 learning library. In *Advances in Neural Information Processing Systems* 32, 2019.

617

618 Leonardo Peroni and Sergey Gorinsky. An end-to-end pipeline perspective on video streaming in
 best-effort networks: A survey and tutorial. *ACM Comput. Surv.*, 2025.

619

620 Jingang Qu, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. TabICL: A tabular
 621 foundation model for in-context learning on large data. In *International Conference on Machine
 622 Learning*, 2025.

623

624 Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*.
 MIT Press, 2006.

625

626 Alexandre A Schoepfer, Jan Weinreich, Ruben Laplaza, Jerome Wasser, and Clemence Corminboeuf.
 Cost-informed bayesian reaction optimization. *Digital Discovery*, 2024.

627

628 Lei Song, Chenxiao Gao, Ke Xue, Chenyang Wu, Dong Li, Jianye Hao, Zongzhang Zhang, and Chao
 629 Qian. Reinforced in-context black-box optimization. *arXiv preprint arXiv:2402.17423*, 2024.

630

631 Richard S Sutton, Andrew G Barto, et al. *Reinforcement Learning: An Introduction*, volume 1. MIT
 Press Cambridge, 1998.

632

633 Shinya Suzuki, Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Multi-
 634 objective bayesian optimization using pareto-frontier entropy. In *International conference on
 635 machine learning*, 2020.

636

637 Jayden Teoh, Pradeep Varakantham, and Peter Vamplew. On generalization across environments
 in multi-objective reinforcement learning, 2025. URL <https://arxiv.org/abs/2503.00799>.

638

639 Michael Volpp, Lukas P Fröhlich, Kirsten Fischer, Andreas Doerr, Stefan Falkner, Frank Hutter,
 640 and Christian Daniel. Meta-learning Acquisition Functions for Transfer Learning in Bayesian
 641 Optimization. *International Conference on Learning Representations*, 2020.

642

643 Boqian Wang, Jiacheng Cai, Chuangui Liu, Jian Yang, and Xianting Ding. Harnessing a novel
 644 machine-learning-assisted evolutionary algorithm to co-optimize three characteristics of an electro-
 645 spun oil sorbent. *ACS Applied Materials & Interfaces*, 12(38):42842–42849, 2020.

646

647 Shukuan Wang, Ke Xue, Song Lei, Xiaobin Huang, and Chao Qian. Monte carlo tree search based
 648 space transfer for black box optimization. *Advances in Neural Information Processing Systems*,
 2024.

648 Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement
649 learning. *Machine learning*, 8(3):229–256, 1992.
650

651 Adam X Yang, Laurence Aitchison, and Henry B Moss. Mongoose: Path-wise smooth bayesian
652 optimisation via meta-learning. *arXiv preprint arXiv:2302.11533*, 2023.
653

654 Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective
655 reinforcement learning and policy adaptation. *Advances in Neural Information Processing Systems*,
656 32, 2019.
657

658 Xinyu Zhang, Daolang Huang, Samuel Kaski, and Julien Martinelli. PABBO: Preferential Amortized
659 Black-Box Optimization. *International Conference on Learning Representations*, 2025.
660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 703 Appendix

704 The appendix is organized as follows:

- 706 • Section A describes the generative process leading to our synthetic pretraining dataset.
- 707 • Section B provides additional details regarding TAMO’s workflow, with Algorithm S1
708 described pretraining.
- 709 • Section C provides a detailed architecture figure along with attention masks for TAMO.
- 710 • Section D provides further details regarding experiments, including model hyperparameters (Section D.3), computational resources (Section D.1) and test functions description (Section D.2).
- 711 • Section E contains additional experiments and analyses:
 - 712 – Single-objective Bayesian optimization
 - 713 – Ablation study on the query set size
 - 714 – Effect of prediction warm-up and the prediction term weight λ_p in the policy training loss
 - 715 – Ablation study on the discount factor γ in the RL objective
 - 716 – Comparison between the standard multi-step and myopic TAMO variants
 - 717 – Effect of model size on optimization performance
 - 718 – Timing breakdown for GP baselines
 - 719 – Evaluation on real-world HPO-3DGS hyperparameter optimization tasks
 - 720 – Effect of pre-training dataset composition
- 721 • Section F displays several examples of GP samples used during pretraining (Section F.1), ex-
722 amples of mean prediction and proposed queries on GP samples (Section F.2, and Section F.3
723 in the decoupled setting).
- 724 • Section G describes to what extent Large Language Models (LLMs) were utilized throughout
725 this work and manuscript.

733 A GENERATIVE PROCESS OF SYNTHETIC PRETRAINING DATASET

735 The model evaluated in Section 5 was trained on a dataset of GP draws. This dataset was constructed
736 to include a variety of configurations, spanning diverse dimensionalities and function properties. All
737 functions were generated inside $[-5.0, 5.0]^{d_x}$ using the following procedure:

- 739 • Input dimensionality $d_x \sim \mathcal{U}(\{1, 2\})$ and output dimensionality $d_y \sim \mathcal{U}(\{1, 2, 3\})$.
- 740 • Regarding output correlations, with probability 1/2, either independent output dimensions
741 are sampled, or they are drawn from a multi-task GP, with task covariance defined as
742 $k(i, j) = (\mathbf{B}\mathbf{B}^T + \text{diag}(\mathbf{v}))_{i,j}, i, j \in \{1, \dots, d_y\}$. In this case, \mathbf{B} is a low-rank matrix
743 with rank $r \sim \mathcal{U}(\{1, \dots, d_y\})$.
- 744 • The data kernel along each output dimension is equally sampled from the RBF, Matérn-
745 3/2, Matérn-5/2 kernels, with standard deviation $\sigma \sim U[0.1, 1.0]$ and lengthscales $l \sim$
746 $\mathcal{N}(2/3, 0.5)$ truncated to the range $[0.1, 2.0]$.
- 747 • The sampled function values, \mathbf{y} , were centered and normalized to lie within $[-1, 1]^{d_y}$.

749 Examples with different dimensionalities are illustrated in Figures S12, S13, S14 and S15 (Sec-
750 tion F.1).

752 B PRETRAINING AND TEST-TIME ALGORITHMS

754 Algorithm S1 and Algorithm S2 describe the pre-training loops and test-time optimization procedure,
755 respectively.

756
757
758
759

760 **Algorithm S1** TAMO Pre-Training Algorithm

761

762 **Require:** task distribution $p(\tau)$, prediction context size N_c , prediction target size N_p , query budget
763 T , number of burn-in iterations η , number of total iterations num_total_iterations

764 1: **for** iteration $i = 1, \dots, \text{num_total_iterations}$ **do** ▷ Prediction task

765 2:

766 3: Sample a task $\tau \sim p(\tau)$

767 4: Sample prediction batches $\mathcal{D}^c = \{(\mathbf{x}_i^c, y_i^c)\}_{i=1}^{N_c}$ and $\mathcal{D}^p = \{\mathbf{x}_i^p\}_{i=1}^{N_p}$ from τ

768 5: Model predicts outcomes: $p(y_{i,k}^p \mid \mathbf{x}_i^p, \mathcal{D}^c), \forall \mathbf{x}_i^p \in \mathcal{D}^p$

769 6: **if** $i \leq \eta$ **then**

770 7: Update model by minimizing the prediction loss $\mathcal{L}^{(p)}$ (Equation 5)

771 8: **else** ▷ Policy learning task after burn-in phase

772 9: Sample a new task $\tau \sim p(\tau)$

773 10: Sample query set \mathcal{D}^q

774 11: Initialize a history set $\mathcal{D}^h \leftarrow \{(\mathbf{x}_0^h, y_0^h)\}, \mathbf{x}_0^h \sim \mathcal{D}^q$

775 12: Set reference point \mathbf{r} and calculate optimal Hypervolume: $\text{HV}^* \leftarrow \text{HV}(\mathcal{P}(\mathcal{X}) \mid \mathbf{r})$

776 13: Initialize Pareto set $\mathcal{P} \leftarrow \{y_0^h\}$

777 14: **for** $t = 1, \dots, T$ **do**

778 15: Select next query point: $\mathbf{x}_t \sim \pi_\theta(\cdot \mid \mathcal{D}^h, t, T)$

779 16: $y_t \leftarrow \text{Evaluate}(\mathbf{x}_t, \tau)$

780 17: Update history set: $\mathcal{D}^h \leftarrow \mathcal{D}^h \cup \{(\mathbf{x}_t, y_t)\}$

781 18: Update Pareto set: $\mathcal{P} \leftarrow \mathcal{P} \cup \{y_t\}$

782 19: Compute reward: $r_t = \frac{\text{HV}(\mathcal{P} \mid \mathbf{r})}{\text{HV}^*}$

783 20: **end for**

784 21: Update model using the overall objective \mathcal{L} (Equation 6)

785 22: **end if**

786 23: **end for**

787
788
789
790
791
792
793

794 **Algorithm S2** TAMO Test-Time Algorithm

795

796 **Require:** Pre-trained TAMO model, new task τ_{test} , query budget T , initial history set $\mathcal{D}_0^h := \{\mathbf{x}^h, y^h\}$
797 (with random samples if empty),

798 1: $\mathcal{D}^h \leftarrow \mathcal{D}_0^h$ ▷ Initialize the history set

799 2: $\mathcal{P} \leftarrow \{y^h\}$ ▷ Initialize the Pareto set

800 3: **for** $t = 1, \dots, T$ **do** ▷ Sample the next query location

801 4: $\mathbf{x}_t \sim \pi_\theta(\cdot \mid \mathcal{D}^h, t, T)$

802 5: $y_t \leftarrow \text{Evaluate}(\mathbf{x}_t, \tau_{\text{test}})$

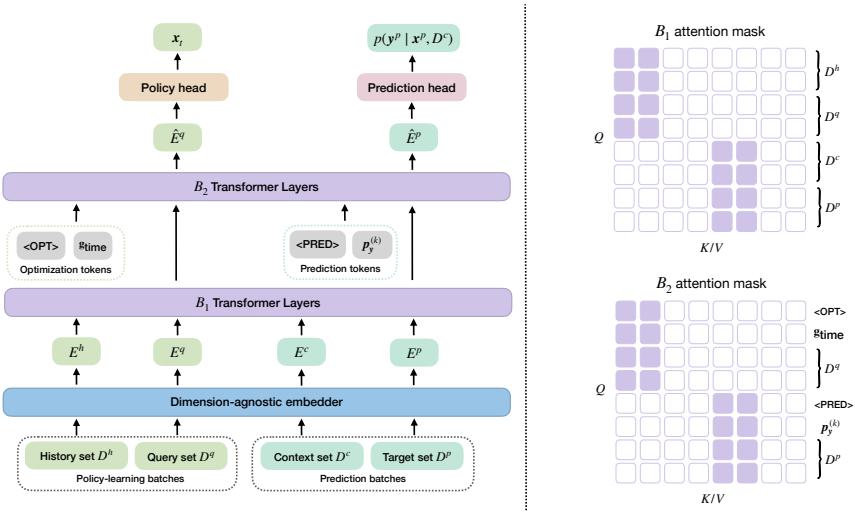
803 6: $\mathcal{D}^h \leftarrow \mathcal{D}^h \cup \{(\mathbf{x}_t, y_t)\}$ ▷ Update the history set

804 7: $\mathcal{P} \leftarrow \mathcal{P} \cup \{y_t\}$ ▷ Update the Pareto set with the new observation

805 8: **end for**

806 9: **return** $\mathcal{D}^h, \mathcal{P}$

807
808
809

810 C DETAILED MODEL ARCHITECTURE
811
812
813832 Figure S1: Detailed Architecture of TAMO.
833
834835 D EXPERIMENTAL DETAILS
836

837 D.1 COMPUTATIONAL RESOURCES

838 We trained TAMO on one NVIDIA H100 80GB HBM3 GPU. All models are evaluated on Tesla
839 V100-SXM2-32GB GPUs.

841 D.2 TEST FUNCTIONS

843 **GP samples optimization.** This benchmark comprises 30 independent multi-output GP draws with
844 $d_x = 2$ inputs and $d_y = 2$ objectives in the *dimensional in-distribution* setting (Section 5.1), and
845 $d_x = 3$, $d_y \in \{2, 3\}$ in the *dimensional out-of-distribution* setting (Section 5.2). We sample each
846 task using the same data-generating process described in Section D.3 and report average performance
847 over the 30 draws.

849 **Ackley–Rastrigin** $d_x = 2, d_y = 2$. Two-objective problem formed by pairing Ackley and Rastrigin
850 and *maximizing* their negations:

$$852 \text{Ackley}(\mathbf{x}) = -20 \exp \left(-0.2 \sqrt{\frac{1}{2} \sum_{i=1}^2 x_i^2} \right) - \exp \left(\frac{1}{2} \sum_{i=1}^2 \cos(2\pi x_i) \right) + e + 20,$$

$$855 \text{Rastrigin}(\mathbf{x}) = 20 + \sum_{i=1}^2 (x_i^2 - 10 \cos(2\pi x_i)),$$

858 And we set $f^{(1)}(\mathbf{x}) = -\text{Ackley}(\mathbf{x})$, $f^{(2)}(\mathbf{x}) = -\text{Rastrigin}(\mathbf{x})$.

860 **Ackley–Rosenbrock** $d_x = 2, d_y = 2$. We pair Ackley (above) with Rosenbrock:

$$862 \text{Rosenbrock}(\mathbf{x}) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2.$$

863 And we set $f^{(1)}(\mathbf{x}) = -\text{Ackley}(\mathbf{x})$, $f^{(2)}(\mathbf{x}) = -\text{Rosenbrock}(\mathbf{x})$.

864
865**Branin–Currin** $d_x = 2, d_y = 2$. Branin:866
867
868869
870
871

$$\text{Branin}(x_1, x_2) = (x_2 - bx_1^2 + cx_1 - r)^2 + s(1 - t) \cos(x_1) + s,$$

where $b = \frac{5.1}{4\pi^2}$, $c = \frac{5}{\pi}$, $r = 6$, $s = 10$, $t = \frac{1}{8\pi}$,

872
873
874875
876

and Currin:

877
878
879
880881
882
883

$$\text{Currin}(z) = \left(1 - e^{-1/(2z_2)}\right) \frac{2300z_1^3 + 1900z_1^2 + 2092z_1 + 60}{100z_1^3 + 500z_1^2 + 4z_1 + 20}.$$

884
885
886
887
888889 We maximize $f^{(1)}(x) = -\text{Branin}(x)$ and $f^{(2)}(x) = -\text{Currin}(x)$.
890891
892
893
894**Oil sorbent** $d_x = 2, d_y = 2$. We also evaluate on the oil-sorbent multi-objective problem (Wang et al., 2020; Daulton et al., 2022). The original task controls a material’s composition and manufacturing with 5 ordinal and 2 continuous parameters to jointly maximize three objectives: oil absorbing capacity, mechanical strength, and water contact angle. In our study, we fix the ordinal parameters to constant values to obtain a 2D continuous design space with the same three objectives.901
902
903
904

905

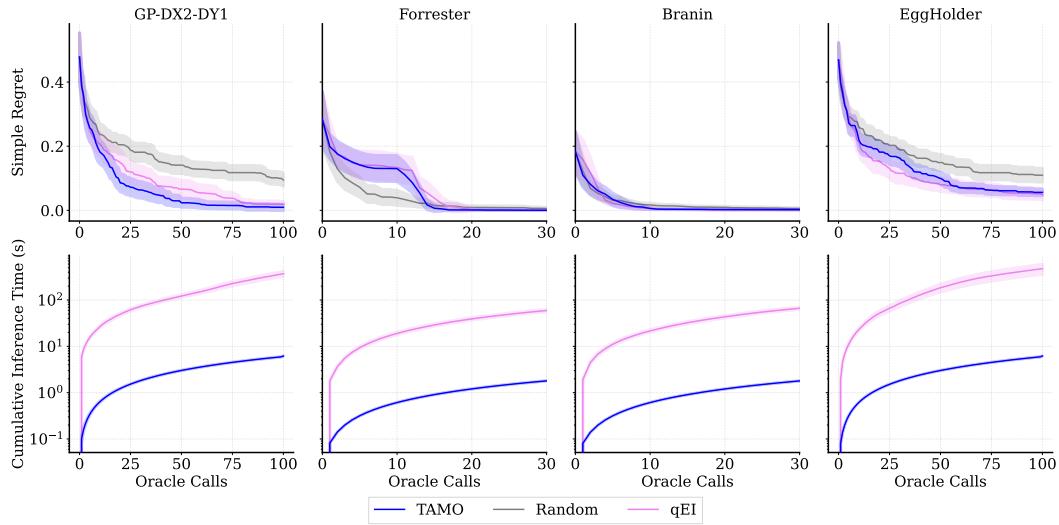
Laser-Plasma $d_x = 4, d_y = 3$. We evaluate on the laser–plasma acceleration dataset (Irshad et al., 2023), which contains 1025 particle-in-cell simulations of a laser wakefield accelerator. Each record provides 4 continuous inputs (plasma density, upramp length, laser focus, downramp length) and 3 objectives (total charge, distance of median, target energy). To obtain a continuous black-box from tabulated simulations, we perform linear interpolation. This task differs in dimensionality from our pretraining distribution, providing an OOD evaluation of cross-dimensional transfer.912
913
914
915
916

917

Normalization. For all problems, we linearly rescale inputs to a common domain $[-5, 5]^{d_x}$ and rescale each objective independently to $[-1, 1]$ prior to logging and hypervolume computation.

918 D.3 HYPERPARAMETERS
919

Dimension-Agnostic Embedder	
Number of learnable positional tokens for x	4
Number of learnable positional tokens for y	3
Number of Transformer layers (L)	4
Dimension of e_x and e_y	64
Transformer Encoder-Decoder	
Dimension of Transformer inputs	64
Point-wise feed-forward dimension of Transformer	256
Number of self-attention layers in Transformer (B)	8
Number of self-attention heads in Transformer	4
Heads	
Number of hidden layers in policy head	3
Number of components in GMM head (K)	20
Number of hidden layers in MLP for each GMM component	3
Training	
Number of iterations	400000
Number of burn-in iterations	393500
Initial learning rate for warm-up iterations (lr_1)	$1 \cdot 10^{-4}$
Initial Learning rate after warm-up (lr_2)	$4 \cdot 10^{-5}$
Learning rate scheduling	Linearly increase from 0 to lr_1 in the first 5% of total iterations; Cosine decay to 0 over total iterations
Size of prediction batch	32
Size of policy-learning batch	16
Weight on prediction loss (λ_{rl})	1.0
discount factor (γ)	1.0
Size of context set	$N_c \sim U[2, 50 \cdot d_x^2]$
Size of target set (N_t)	$300 - N_c$
Size of query set (N_q)	256
Optimization budget T	100
Noise level σ	0.0
Number of initial observations during pretraining	1
Evaluation	
Number of initial observations during test time	1
Noise level σ	0.0
Size of query set (N_q)	2048
Optimization budget (T)	100

943 Table S1: Hyperparameter settings for TAMO evaluated in Section 5.
944945 E ADDITIONAL EXPERIMENTS
946947 Figure S2: Simple regret and inference time on **synthetic examples for single-objective optimization**.
948 Mean with 95% confidence intervals computed across 30 runs with random initial observations.
949 Again, TAMO matches state-of-the-art regret while dramatically reducing proposal time.
950

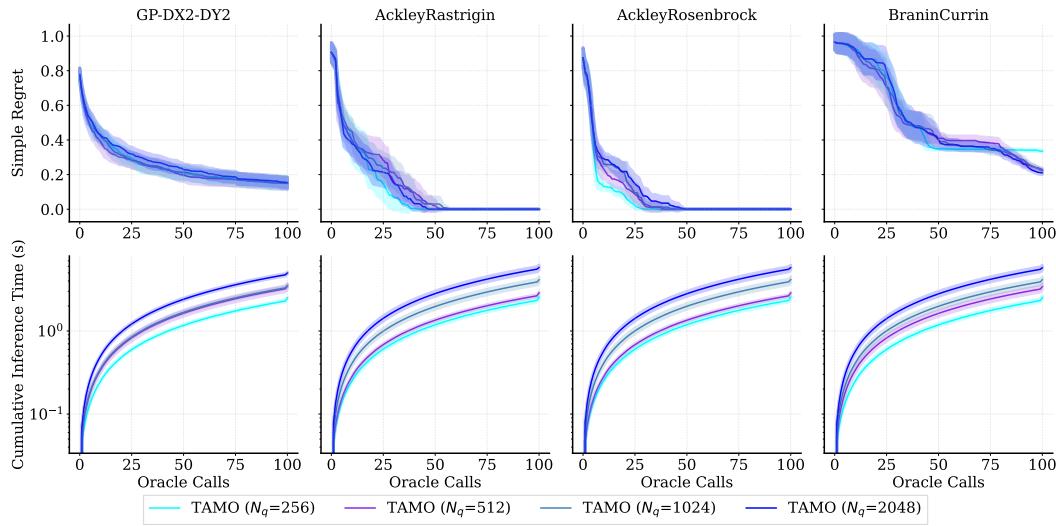


Figure S3: **Effect of query set size.** Simple regret (top) and cumulative inference time (bottom) for TAMO with $N_q \in \{256, 512, 1024, 2048\}$ on four synthetic tasks. Means with 95% CIs over 30 runs.

Larger N_q increases wall-clock roughly linearly while leaving regret essentially unchanged.

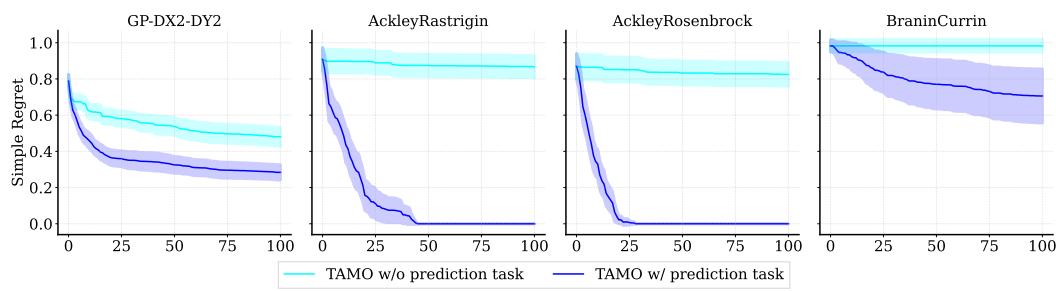


Figure S4: **Removing the prediction warm-up and prediction term from the training loss (Equation 6).** Simple regret on four synthetic tasks. Means with 95% CIs over 30 runs. **Introducing an auxiliary prediction task before and during policy training is decisive.**

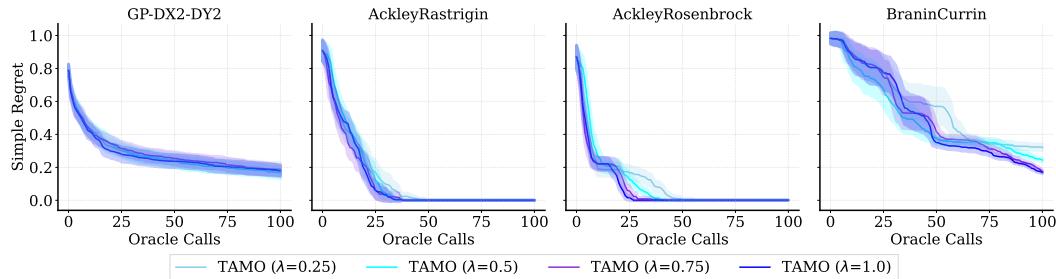


Figure S5: **Effect of the prediction term weight λ_p in the training loss (Equation 6) during policy training.** Simple regret on four synthetic tasks. Means with 95% CIs over 30 runs. **Once policy training starts, performance is relatively insensitive to λ_p , with slightly better results for larger weights.**

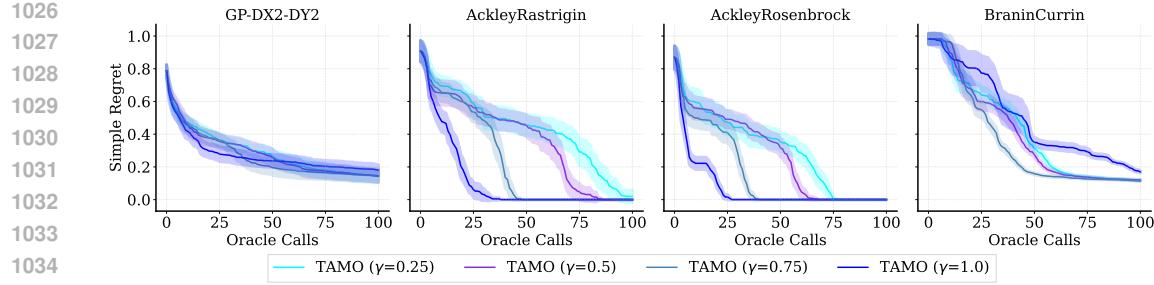


Figure S6: **Effect of the discount factor γ in the RL objective (Equation 4) during policy training.** Simple regret on four synthetic tasks. Means with 95% CIs over 30 runs. Overall, performance is fairly robust to γ , with all policies eventually reaching similar regret levels, and a general trend toward faster early regret reduction as the degree of non-myopicity (larger γ) increases.

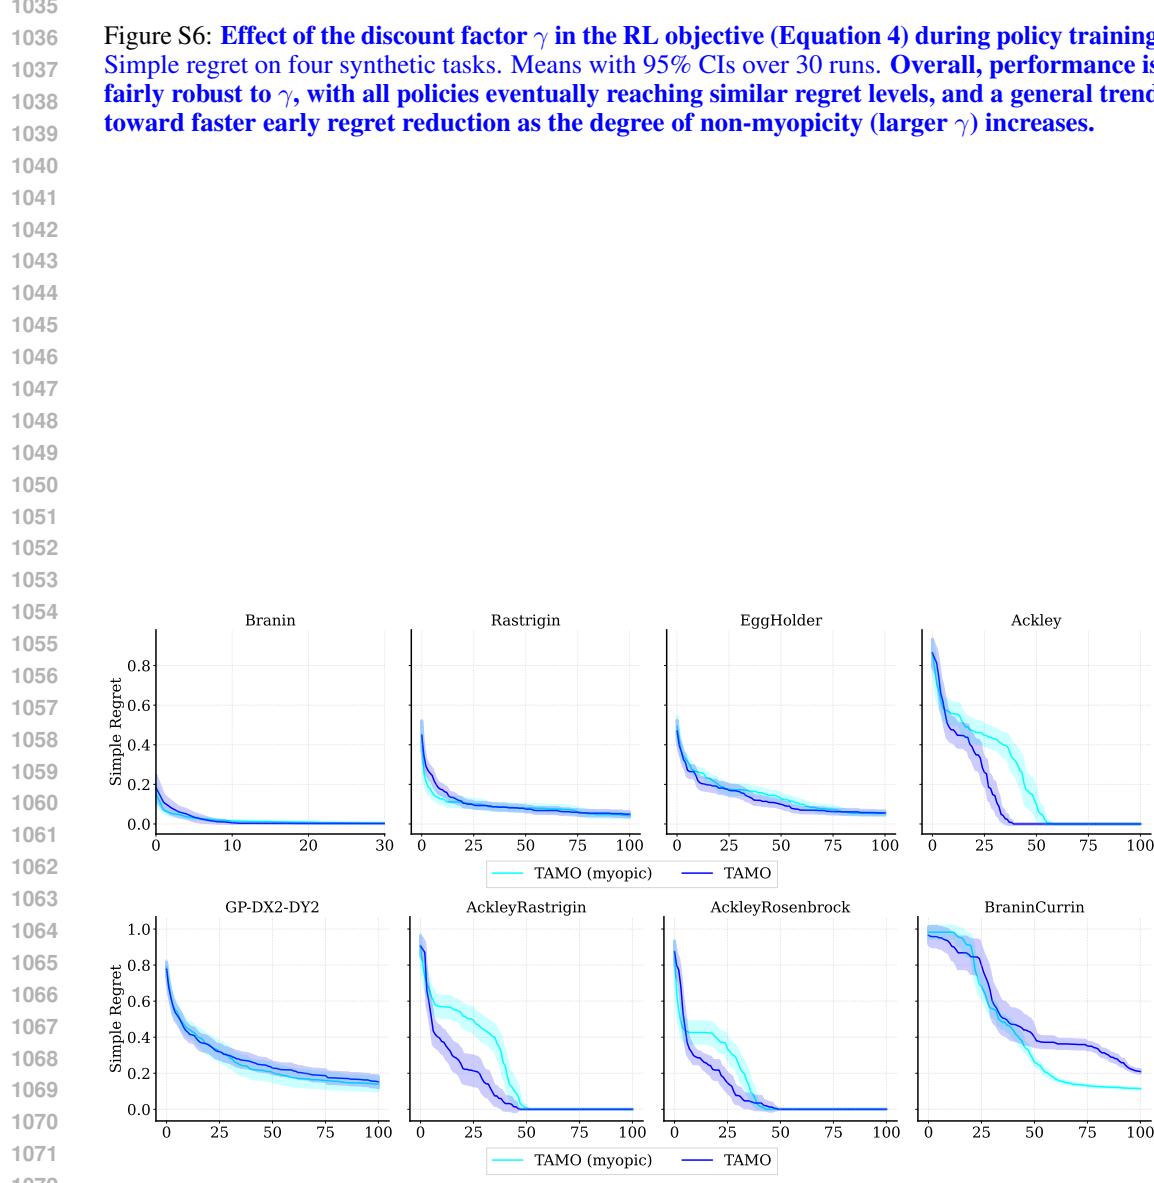


Figure S7: **Effect of the pretraining horizon in the RL objective (Equation 4) during policy training.** We compare a myopic variant of TAMO, pretrained with horizon $T = 1$, to the standard TAMO pretrained with $T = 100$. For the non-myopic model, each pretraining episode starts from a single randomly sampled context point, and the policy acts for up to 100 steps. For the myopic case, each episode starts from a randomly sampled context set, and the policy proposes a single additional point, so every decision is strictly one-step. Simple regret on four synthetic tasks for single objective optimization (top) and multi-objective optimization (bottom). Means with 95% CIs over 30 runs. Results clearly advocate for longer pretraining horizons, except for the BraninCurrin problem.

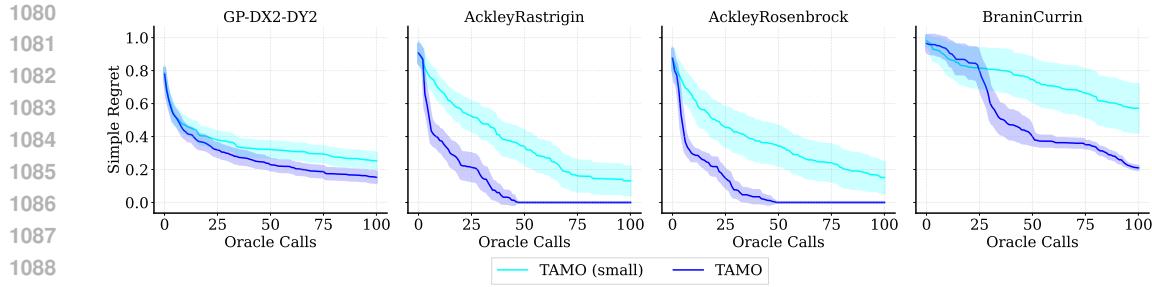


Figure S8: **Effect of model size on optimization performance.** We compare a smaller TAMO variant, trained for the same number of iterations as the standard model but using 2 Transformer layers per module (dimension-agnostic embedder, B_1 layers, B_2 layers) instead of 4. Simple regret on four synthetic tasks. Means with 95% CIs over 30 runs. **While the smaller model remains competitive, the larger backbone consistently attains lower regret, especially on the more challenging tasks.**

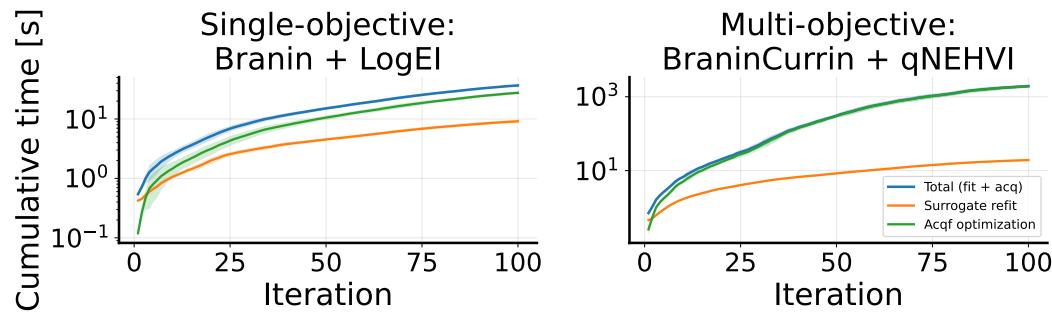
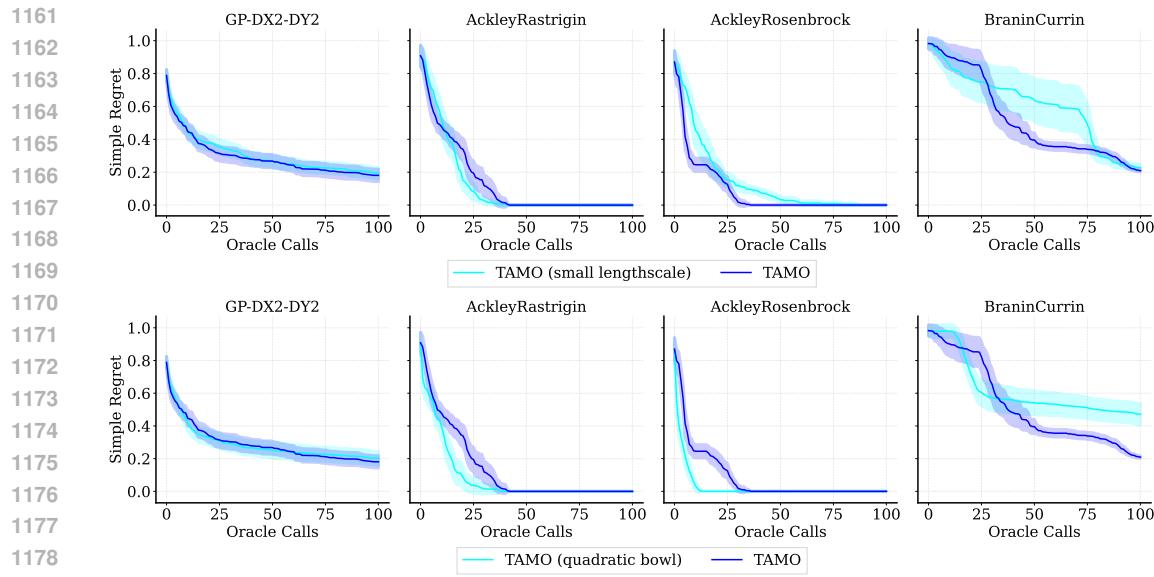
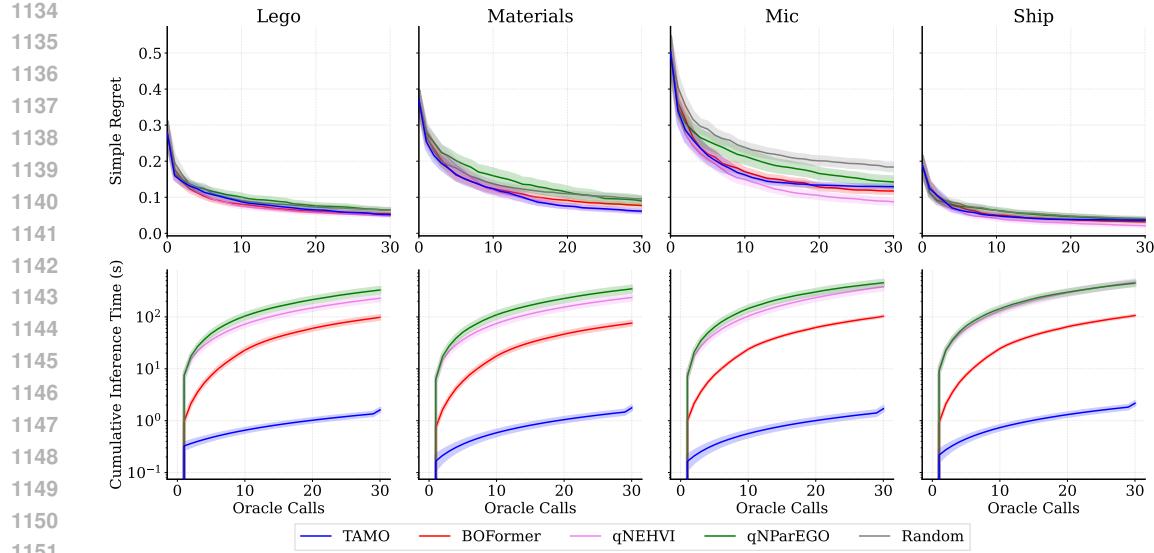
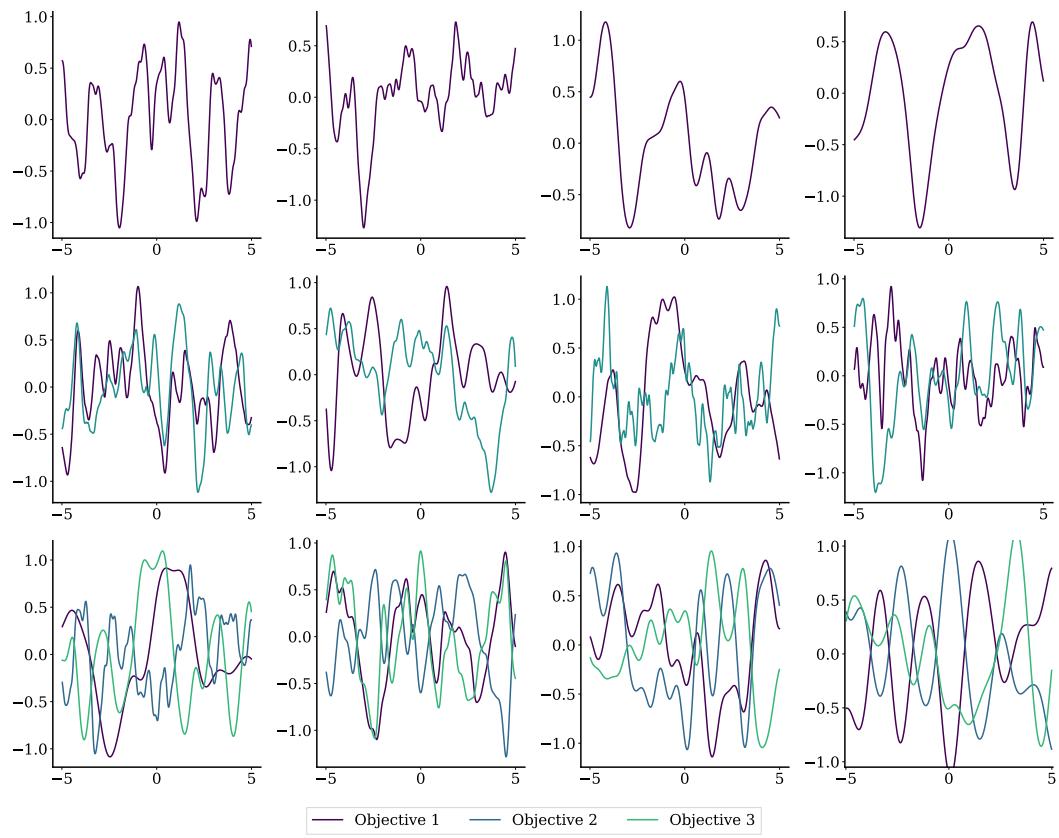
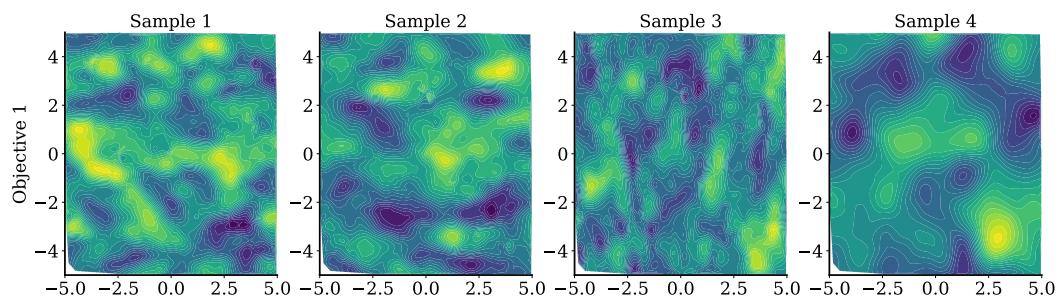


Figure S9: **Timing breakdown for GP baselines.** Left: single-objective Branin with EI. Right: multi-objective Branin–Currin with qNEHVI. Curves show cumulative inference time (mean \pm 1 std over 10 runs), decomposed into total time (blue), surrogate refits (orange), and acquisition-function optimization (green). **Acquisition optimization quickly dominates the cost in both cases, and especially in the multi-objective setting.**



1188 F VISUALIZATION EXAMPLES
11891190 F.1 EXAMPLES OF GP SAMPLES FROM PRETRAINING
11911192 Figures S12 to S15 show some examples of the GP samples used for pre-training.
11931221 Figure S12: GP Samples used during pretraining with input dimension $d_x = 1$ and output dimension
1222 $d_y = 1, 2, 3$.
12231234 Figure S13: GP Samples used during pretraining with input dimension $d_x = 2$ and output dimension
1235 $d_y = 1$.
12361238 F.2 EXAMPLES OF INFERENCE
12391240 Figures S16 and Figure S17 show examples of mean predictions and proposed queries within a total
1241 budget $T = 100$ on GP samples, for input dimension $d_x = 2$ and output dimensions $d_y = 1$ and
 $d_y = 2$, respectively.
1242

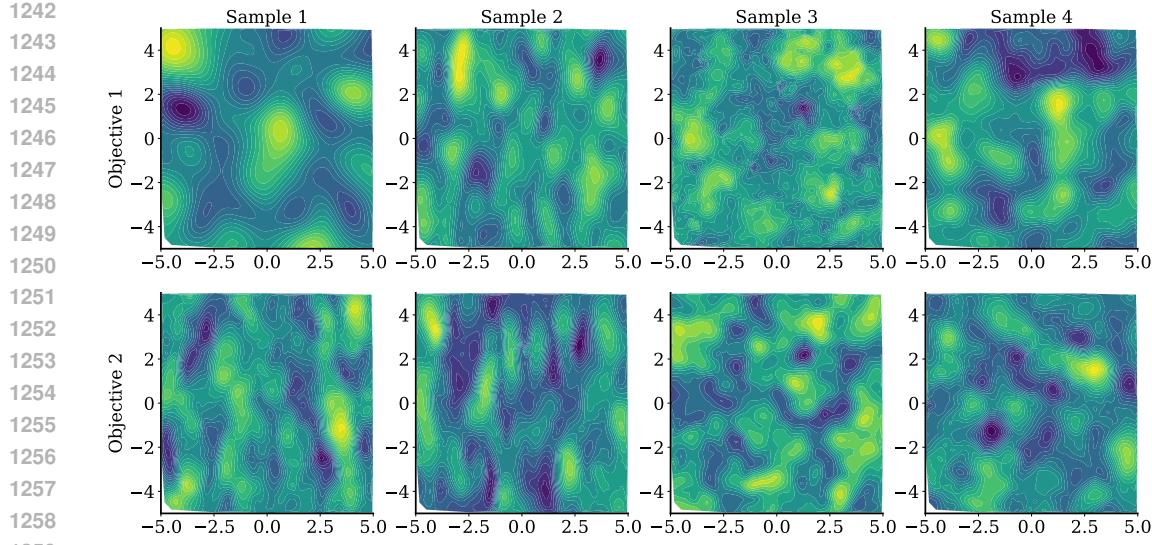


Figure S14: GP Samples used during pretraining with input dimension $d_x = 2$ and output dimension $d_y = 2$.

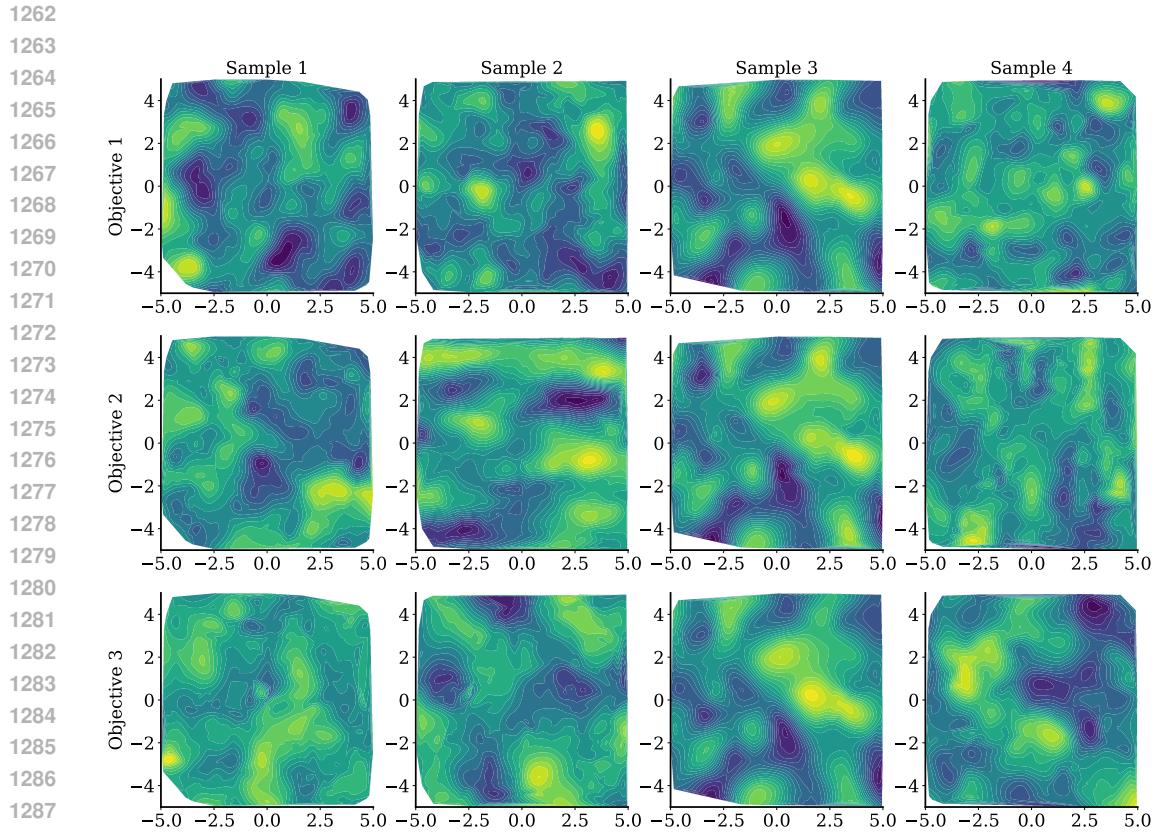
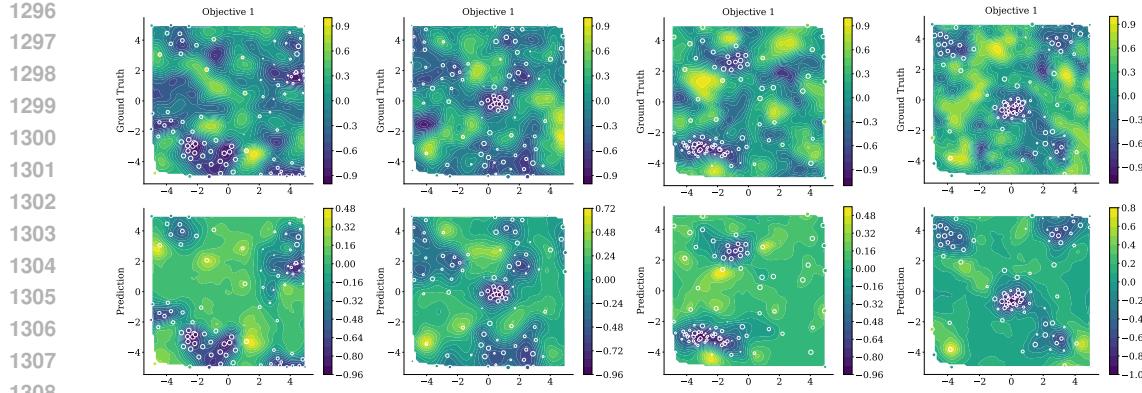


Figure S15: GP Samples used during pretraining with input dimension $d_x = 2$ and output dimension $d_y = 3$.

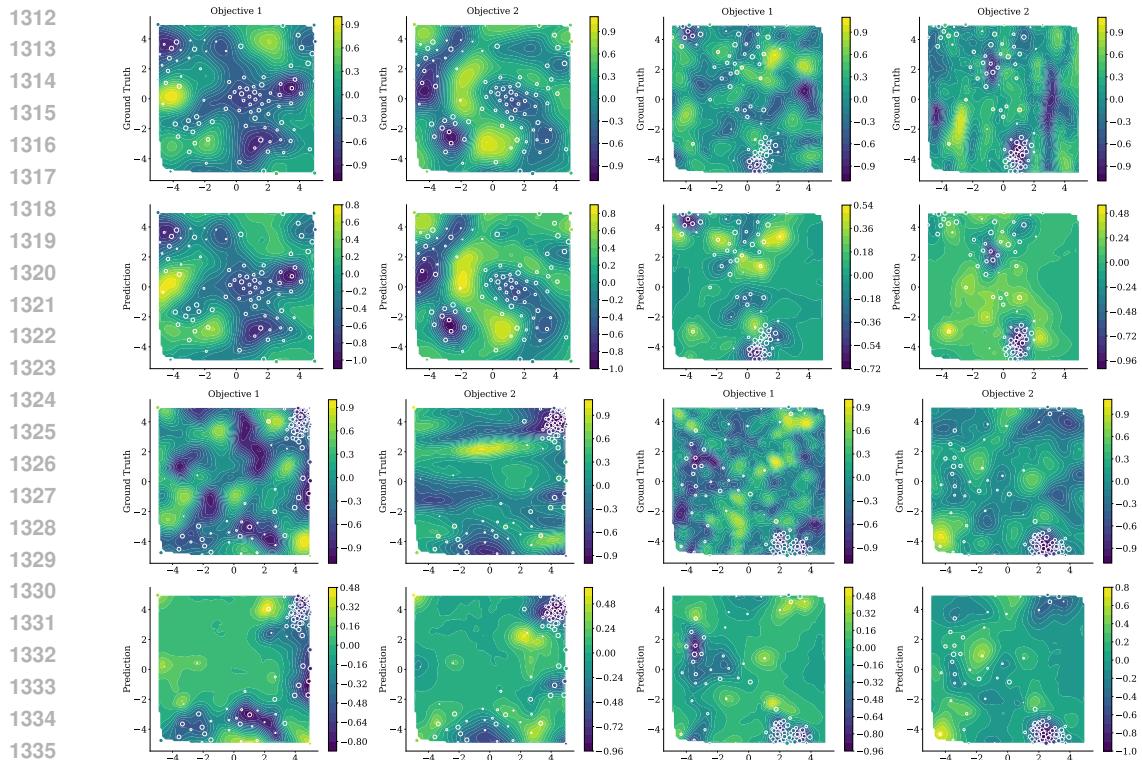
1291
1292
1293
1294
1295

F.3 EXAMPLES OF DECOUPLED OBSERVATIONS

Figure S18 shows examples of mean predictions and proposed queries under the decoupled setting.



1309
1310
Figure S16: Inference on GP examples ($d_x = 2, d_y = 1$), with query points proposed over 100
1311
optimization steps (white circle, size increasing along with the number of queries).



1337
1338
Figure S17: Inference on GP examples ($d_x = 2, d_y = 2$), with query points proposed over 100
1339
optimization steps (white circles, size increasing along with the number of queries).

G USE OF LARGE LANGUAGE MODELS

1342
1343
We employed LLMs to support the following aspects of our research:

- 1344
1345
• Ideation & Exploration: Assisting with brainstorming of methods and conducting preliminary literature searches and summarizations.
- 1346
1347
• Coding Assistance: Generating boilerplate PyTorch code, visualization scripts, and test structures. All LLM-generated code was manually reviewed and verified by the authors.
- 1348
1349
• Writing Assistance: Refining sentence structure, grammar, and clarity in the manuscript. The scientific content and all claims remain the sole work of the authors.

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

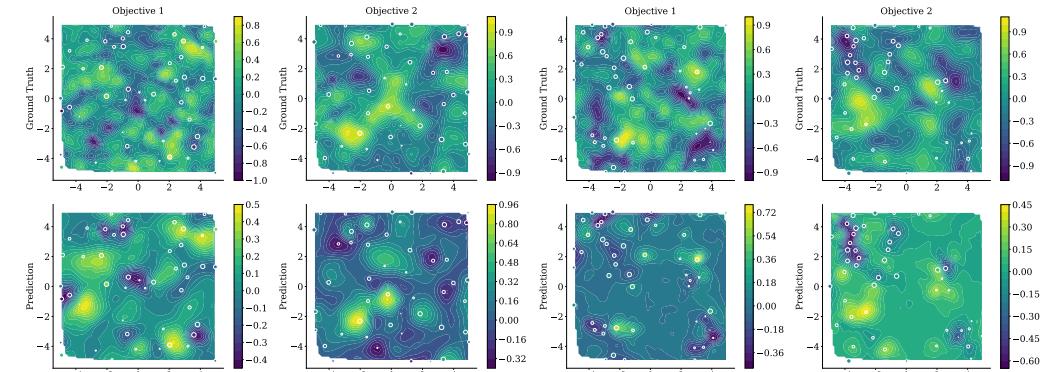


Figure S18: Mean predictions and queries on GP examples ($d_x = 2, d_y = 2$) from TAMO under the decoupled setting. Each column represents a distinct objective; queries to evaluate that objective are outlined by circles, with the sizes increasing over time to show the optimization progress. Note no queries overlap between objectives.

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403