
Multi-Fidelity Active Learning with GFlowNets

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In the last decades, the capacity to generate large amounts of data in science and
2 engineering applications has been growing steadily. Meanwhile, the progress in
3 machine learning has turned it into a suitable tool to process and utilise the available
4 data. Nonetheless, many relevant scientific and engineering problems present
5 challenges where current machine learning methods cannot yet efficiently leverage
6 the available data and resources. For example, in scientific discovery, we are often
7 faced with the problem of exploring very large, high-dimensional spaces, where
8 querying a high fidelity, black-box objective function is very expensive. Progress
9 in machine learning methods that can efficiently tackle such problems would help
10 accelerate currently crucial areas such as drug and materials discovery. In this paper,
11 we propose the use of GFlowNets for multi-fidelity active learning, where multiple
12 approximations of the black-box function are available at lower fidelity and cost.
13 GFlowNets are recently proposed methods for amortised probabilistic inference
14 that have proven efficient for exploring large, high-dimensional spaces and can
15 hence be practical in the multi-fidelity setting too. Here, we describe our algorithm
16 for multi-fidelity active learning with GFlowNets and evaluate its performance in
17 both well-studied synthetic tasks and practically relevant applications of molecular
18 discovery. Our results show that multi-fidelity active learning with GFlowNets
19 can efficiently leverage the availability of multiple oracles with different costs and
20 fidelities to accelerate scientific discovery and engineering design.

21 1 Introduction

22 The current most pressing challenges for humanity, such as the climate crisis and the threat of
23 pandemics or antibiotic resistance could be tackled, at least in part, with new scientific discoveries.
24 By way of illustration, materials discovery can play an important role in improving the energy
25 efficiency of energy production and storage; and reducing the costs and duration for drug discovery
26 has the potential to more effectively and rapidly mitigate the consequences of new diseases. In
27 recent years, researchers in materials science, biochemistry and other fields have increasingly adopted
28 machine learning as a tool as it holds the promise to drastically accelerate scientific discovery
29 [7, 67, 3, 12].

30 Although machine learning has already made a positive impact in scientific discovery applications
31 [55, 27], unleashing its full potential will require improving the current algorithms [1]. For example,
32 typical tasks in potentially impactful applications in materials and drug discovery require exploring
33 combinatorially large, high-dimensional spaces [46, 6], where only small, noisy data sets are available,
34 and obtaining new annotations computationally or experimentally is very expensive. Such scenarios
35 present serious challenges even for the most advanced current machine learning methods.

36 In the search for a useful discovery, we typically define a quantitative proxy for usefulness, which we
37 can view as a black-box function. One promising avenue for improvement is developing methods
38 that more efficiently leverage the availability of multiple approximations of the target black-box

39 function at lower fidelity but much lower cost than the highest fidelity oracle [10, 14]. For example,
40 the most accurate estimation of the properties of materials and molecules is only typically obtained
41 via synthesis and characterisation in a laboratory. However, this is only feasible for a small number of
42 promising candidates. Approximate quantum mechanics simulations of a larger amount of chemical
43 compounds can be performed via Density Functional Theory (DFT) [41, 51]. However, DFT is
44 still computationally too expensive for high-throughput exploration of large search spaces. Thus,
45 large-scale exploration can only be achieved through cheaper but less accurate oracles. Nonetheless,
46 solely relying on low-fidelity approximations is clearly suboptimal. Ideally, such tasks would be best
47 tackled by methods that can efficiently and adaptively distribute the available computational budget
48 between the multiple oracles depending on the already acquired information.

49 The past decade has seen significant progress in multi-fidelity Bayesian optimisation (BO) [19, 53],
50 including methods that leverage the potential of deep neural networks [36]. Although highly relevant
51 for scientific discovery, standard BO is not perfectly suited for some of the challenges in materials and
52 drug discovery tasks. First and foremost, BO’s ultimate goal is to find the optimum of an expensive
53 black-box function. However, even the highest fidelity oracles in such problems are underspecified
54 with respect to the actual, relevant, downstream applications. Therefore, it is imperative to develop
55 methods that, instead of “simply” finding the optimum, discover a set of diverse, high-scoring
56 candidates.

57 Recently, generative flow networks (GFlowNets) [4] have demonstrated their capacity to find diverse
58 candidates through discrete probabilistic modelling, with particularly promising results when embed-
59 ded in an active learning loop [22]. Here, we propose to extend the applicability of GFlowNets for
60 multi-fidelity active learning.

61 In this paper, we present an algorithm for multi-fidelity active learning with GFlowNets. We provide
62 empirical results in two synthetic benchmark tasks and four practically relevant tasks for biological
63 sequence design and molecular modelling. As a main result, we demonstrate that multi-fidelity
64 active learning with GFlowNets discovers diverse, high-scoring samples when multiple oracles with
65 different fidelities and costs are available, with lower computational cost than its single-fidelity
66 counterpart.

67 2 Related Work

68 Our work can be framed within the broad field of active learning (AL), a class of machine learning
69 methods whose goal is to learn an efficient data sampling scheme to accelerate training [50]. For the
70 bulk of the literature in AL, the goal is to train an accurate model $h(x)$ of an unknown target function
71 $f(x)$, as in classical supervised learning. However, in certain scientific discovery problems, which is
72 the motivation of our work, a desirable goal is often to discover multiple, diverse candidates x with
73 high values of $f(x)$. The reason is that the ultimate usefulness of a discovery is extremely expensive
74 to quantify and we always rely on more or less accurate approximations. Since we generally have
75 the option to consider more than one candidate solution, it is safer to generate a set of diverse and
76 apparently good solutions, instead of focusing on the single global optimum of the wrong function.

77 This distinctive goal is closely connected to related research areas such as Bayesian optimisation [19]
78 and active search [20]. Bayesian optimisation (BO) is an approach grounded in Bayesian inference
79 for the problem of optimising a black-box objective function $f(x)$ that is expensive to evaluate. In
80 contrast to the problem we address in this paper, standard BO typically considers continuous domains
81 and works best in relatively low-dimensional spaces [18]. Nonetheless, in recent years, approaches
82 for BO with structured data [13] and high-dimensional domains [21] have been proposed in the
83 literature. The main difference between BO and the problem we tackle in this paper is that we are
84 interested in finding multiple, diverse samples with high value of f and not only the optimum.

85 This goal, as well as the discrete nature of the search space, is shared with active search, a variant of
86 active learning in which the task is to efficiently find multiple samples of a valuable (binary) class
87 from a discrete domain \mathcal{X} [20]. This objective was already considered in the early 2000s by Warmuth
88 et al. for drug discovery [59], and more formally analysed in later work [26, 25]. A recent branch of
89 research in stochastic optimisation that considers diversity is so-called Quality-Diversity [9], which
90 typically uses evolutionary algorithms that perform search in the latent space. All these and other
91 problems such as multi-armed bandits [48] and the general framework of experimental design [8] all

92 share the objective of optimising or exploring an expensive black-box function. Formal connections
 93 between some of these areas have been established in the literature [54, 17, 23, 15].

94 Multi-fidelity methods have been proposed in most of these related areas of research. An early
 95 survey on multi-fidelity methods for Bayesian optimisation was compiled by Peherstorfer et al. [42],
 96 and research on the subject has continued since [44, 53], with the proposal of specific acquisition
 97 functions [56] and the use of deep neural networks to improve the modelling [36]. Interestingly, the
 98 literature on multi-fidelity active learning [35] is scarcer than on Bayesian optimisation. Recently,
 99 works on multi-fidelity active search have also appeared in the literature [40]. Finally, multi-fidelity
 100 methods have recently started to be applied in scientific discovery problems [10, 14]. However, the
 101 literature is still scarce probably because most approaches do not tackle the specific needs in scientific
 102 discovery, such as the need for diverse samples. Here, we aim addressing this need with the use of
 103 GFlowNets [4, 24] for multi-fidelity active learning.

104 3 Method

105 3.1 Background

106 **GFlowNets** Generative Flow Networks [GFlowNets; 4, 5] are amortised samplers designed for
 107 sampling from discrete high-dimensional distributions. Given a space of compositional objects \mathcal{X}
 108 and a non-negative reward function $R(x)$, GFlowNets are designed to learn a stochastic policy π
 109 that generates $x \in \mathcal{X}$ with a probability proportional to the reward, that is $\pi(x) \propto R(x)$. This distinctive
 110 property induces sampling diverse, high-reward objects, which is a desirable property for scientific
 111 discovery, among other applications [23].

112 The objects $x \in \mathcal{X}$ are constructed sequentially by sampling transitions $s_t \rightarrow s_{t+1} \in \mathbb{A}$ between
 113 partially constructed objects (states) $s \in \mathcal{S}$, which includes a unique empty state s_0 . The stochastic
 114 forward policy is typically parameterised by a neural network $P_F(s_{t+1}|s_t; \theta)$, where θ denotes the
 115 learnable parameters, which models the distribution over transitions $s_t \rightarrow s_{t+1}$ from the current state
 116 s_t to the next state s_{t+1} . The backward transitions are parameterised too and denoted $P_B(s_t|s_{t+1}; \theta)$.
 117 The probability $\pi(x)$ of generating an object x is given by P_F and its sequential application:

$$\pi(x) = \sum_{\tau: s_{|\tau|-1} \rightarrow x \in \tau} \prod_{t=0}^{|\tau|-1} P_F(s_{t+1}|s_t; \theta),$$

118 which sums over all trajectories τ with terminating state x , where $\tau = (s_0 \rightarrow s_1 \dots \rightarrow s_{|\tau|})$ is a
 119 complete trajectory. To learn the parameters θ such that $\pi(x) \propto R(x)$ we use the trajectory balance
 120 learning objective [37]

$$\mathcal{L}_{TB}(\tau; \theta) = \left(\log \frac{Z_\theta \prod_{t=0}^n P_F(s_{t+1}|s_t; \theta)}{R(x) \prod_{t=1}^n P_B(s_t|s_{t+1}; \theta)} \right)^2, \quad (1)$$

121 where Z_θ is an approximation of the partition function $\sum_{x \in \mathcal{X}} R(x)$ that is learned. The GFlowNet
 122 learning objective supports training from off-policy trajectories, so for training the trajectories are
 123 typically sampled from a mixture of the current policy with a uniform random policy. The reward is
 124 also tempered to make the policy focus on the modes.

125 **Active Learning** In its simplest formulation, the active learning problem that we consider is as
 126 follows: we start with an initial data set $\mathcal{D} = \{(x_i, f(x_i))\}$ of samples $x \in \mathcal{X}$ and their evaluations
 127 by an expensive, black-box objective function (oracle) $f: \mathcal{X} \rightarrow \mathbb{R}$, which we use to train a surrogate
 128 model $h(x)$. A GFlowNet can then be trained to learn a generative policy $\pi_\theta(x)$ using $h(x)$ as reward
 129 function, that is $R(x) = h(x)$. Optionally, we can instead train a probabilistic proxy $p(f|\mathcal{D})$ and use
 130 as reward the output of an acquisition function $\alpha(x, p(f|\mathcal{D}))$ that considers the epistemic uncertainty
 131 of the surrogate model, as typically done in Bayesian optimisation. Finally, we use the policy $\pi(x)$ to
 132 generate a batch of samples to be evaluated by the oracle f , we add them to our data set and repeat
 133 the process a number of active learning rounds.

134 While much of the active learning literature [50] has focused on so-called *pool-based* active learning,
 135 where the learner selects samples from a pool of unlabelled data, we here consider the scenario of
 136 *de novo query synthesis*, where samples are selected from the entire object space \mathcal{X} . This scenario

137 is particularly suited for scientific discovery [30, 62, 64, 33]. The ultimate goal pursued in active
 138 learning applications is also heterogeneous. Often, the goal is the same as in classical supervised
 139 machine learning: to train an accurate (proxy) model $h(x)$ of the unknown target function $f(x)$. For
 140 some problems in scientific discovery, we are usually not interested in the accuracy in the entire input
 141 space \mathcal{X} , but rather in discovering new, diverse objects with high values of f . This is connected to
 142 other related problems such as Bayesian optimisation [19], active search [20] or experimental design
 143 [8], as reviewed in Section 2.

144 3.2 Multi-Fidelity Active Learning

145 We now consider the following active learning problem with multiple oracles of different fidelities.
 146 Our ultimate goal is to generate a batch of K samples $x \in \mathcal{X}$ according to the following desiderata:

- 147 • The samples obtain a high value when evaluated by the objective function $f : \mathcal{X} \rightarrow \mathbb{R}^+$.
- 148 • The samples in the batch should be distinct and diverse, that is cover distinct high-valued
 149 regions of f .

150 Furthermore, we are constrained by a computational budget Λ that limits our capacity to evaluate f .
 151 While f is extremely expensive to evaluate, we have access to a discrete set of surrogate functions
 152 (oracles) $\{f_m\}_{1 \leq m \leq M} : \mathcal{X} \rightarrow \mathbb{R}^+$, where m represents the fidelity index and each oracle has an
 153 associated cost λ_m . We assume $f_M = f$ because there may be even more accurate oracles for
 154 the true usefulness but we do not have access to them, which means that even when measured by
 155 $f = f_M$, diversity remains an important objective. We also assume, without loss of generality, that
 156 the larger m , the higher the fidelity and that $\lambda_1 < \lambda_2 < \dots < \lambda_M$. This scenario resembles many
 157 practically relevant problems in scientific discovery, where the objective function f_M is indicative but
 158 not a perfect proxy of the true usefulness of objects x —hence we want diversity—yet it is extremely
 159 expensive to evaluate—hence cheaper, approximate models are used in practice.

160 In multi-fidelity active learning—as well as in multi-fidelity Bayesian optimisation—the iterative
 161 sampling scheme consists of not only selecting the next object x (or batch of objects) to evaluate, but
 162 also the level of fidelity m , such that the procedure is cost-effective.

163 Our algorithm, MF-GFN, detailed in Algorithm 1, proceeds as follows: An active learning round j
 164 starts with a data set of annotated samples $\mathcal{D}_j = \{(x_i, f_m(x_i), m_i)\}_{1 \leq m \leq M}$. The data set is used to
 165 fit a probabilistic *multi-fidelity surrogate* model $h(x, m)$ of the posterior $p(f_m(x)|x, m, \mathcal{D})$. We use
 166 Gaussian Processes [47], as is common in Bayesian optimisation, to model the posterior, such that the
 167 model h predicts the conditional Gaussian distribution of $f_m(x)$ given (x, m) and the existing data
 168 set \mathcal{D} . We implement a multi-fidelity GP kernel by combining a Matern kernel evaluated on x with a
 169 linear downsampling kernel over m [61]. In the higher dimensional problems, we use Deep Kernel
 170 Learning [60] to increase the expressivity of the surrogate models. The candidate x is modelled with
 171 the deep kernel while the fidelity m is modelled with the same linear downsampling kernel. The output
 172 of the proxy model is then used to compute the value of a *multi-fidelity acquisition function* $\alpha(x, m)$.
 173 In our experiments, we use the multi-fidelity version [56] of Max-Value Entropy Search (MES) [58],
 174 which is an information-theoretic acquisition function widely used in Bayesian optimization. MES
 175 aims to maximize the mutual information between the value of the queried x and the maximum value
 176 attained by the objective function, f^* . The multi-fidelity variant is designed to select the candidate x
 177 and the fidelity m that maximize the mutual information between f_M^* and the oracle at fidelity m ,
 178 f_m , weighted by the cost of the oracle:

$$\alpha(x, m) = \frac{1}{\lambda_m} I(f_M^*; f_m | \mathcal{D}_j). \quad (2)$$

179 We provide further details about the acquisition function in Appendix A. A multi-fidelity acquisition
 180 function can be regarded as a cost-adjusted utility function. Therefore, in order to carry out a cost-
 181 aware search, we seek to sample diverse objects with high value of the acquisition function. In this
 182 paper, we propose to use a GFlowNet as a generative model trained for this purpose (see further
 183 details below in Section 3.3). An active learning round terminates by generating N objects from
 184 the sampler (here the GFlowNet policy π) and forming a batch with the best B objects, according
 185 to α . Note that $N \gg B$, since sampling from a GFlowNet is relatively inexpensive. The selected

186 objects are annotated by the corresponding oracles and incorporated into the data set, such that
 187 $\mathcal{D}_{j+1} = \mathcal{D}_j \cup \{(x_1, f_m(x_1), m_1), \dots, (x_B, f_m(x_B), m_B)\}$.

Algorithm 1: MF-GFN: Multi-fidelity active learning with GFlowNets. See Section 4.1 for quality (Top- $K(\mathcal{D})$) and diversity metrics.

Input: $\{(f_m, \lambda_m)\}$: M oracles and their corresponding costs;
 $\mathcal{D}_0 = \{(x_i, f_m(x_i), m_i)\}$: Initial dataset;
 $h(x, m)$: Multi-fidelity Gaussian Process proxy model;
 $\alpha(x, m)$: Multi-fidelity acquisition function;
 $R(\alpha(x), \beta)$: reward function to train the GFlowNet;
 B : Batch size of oracles queries;
 Λ : Maximum available budget;
 K : Number of top-scoring candidates to be evaluated at termination;
Result: Top- $K(\mathcal{D})$, Diversity
Initialization: $\Lambda_j = 0, \mathcal{D} = \mathcal{D}_0$
while $\Lambda_j < \Lambda$ **do**
 • Fit h on dataset \mathcal{D} ;
 • Train GFlowNet with reward $R(\alpha(x), \beta)$ to obtain policy $\pi_\theta(x)$;
 • Sample B tuples $(x_i, m_i) \sim \pi_\theta$;
 • Evaluate each tuple with the corresponding oracle to form batch
 $\mathcal{B} = \{(x_1, f_m(x_1), m_1), \dots, (x_B, f_m(x_B), m_B)\}$;
 • Update dataset $\mathcal{D} = \mathcal{D} \cup \mathcal{B}$;
end

188 3.3 Multi-Fidelity GFlowNets

189 In order to use GFlowNets in the multi-fidelity active learning loop described above, we propose to
 190 make the GFlowNet sample the fidelity m for each object $x \in \mathcal{X}$ in addition to x itself. Formally,
 191 given a baseline GFlowNet with state and transition spaces \mathcal{S} and \mathbb{A} , we augment the state space with
 192 a new dimension for the fidelity $\mathcal{M}' = \{0, 1, 2, \dots, M\}$ (including $m = 0$, which corresponds to
 193 unset fidelity), such that the augmented, multi-fidelity space is $\mathcal{S}_{\mathcal{M}'} = \mathcal{S} \cup \mathcal{M}'$. The set of allowed
 194 transitions $\mathbb{A}_{\mathcal{M}'}$ is augmented such that a fidelity $m > 0$ of a trajectory must be selected once, and
 195 only once, from any intermediate state.

196 Intuitively, allowing the selection of the fidelity at any step in the trajectory should give flexibility for
 197 better generalisation. At the end, finished trajectories are the concatenation of an object x and the
 198 fidelity m , that is $(x, m) \in \mathcal{X}_{\mathcal{M}'} = \mathcal{X} \cup \mathcal{M}'$. In summary, the proposed approach enables to jointly
 199 learn the policy that samples objects in a potentially very large, high-dimensional space, together
 200 with the level of fidelity, that maximise a given multi-fidelity acquisition function as reward.

201 4 Empirical Evaluation

202 In this section, we describe the evaluation metrics and experiments performed to assess the validity
 203 and performance of our proposed approach of multi-fidelity active learning with GFlowNets. Overall,
 204 the purpose of this empirical evaluation is to answer the following questions:

- 205 • **Question 1:** Is our multi-fidelity active learning approach able to find high-scoring, diverse
 206 samples at lower cost than active learning with a single oracle?
- 207 • **Question 2:** Does our proposed multi-fidelity GFlowNet, which learns to sample fidelities
 208 together with objects (x, m) , provide any advantage over sampling only objects x ?

209 In Section 4.1 we describe the metrics proposed to evaluate the performance our proposed method,
 210 as well as the baselines, which we describe in Section 4.2. In Section 4.3, we present results on
 211 synthetic tasks typically used in the multi-fidelity BO and active learning literature. In Section 4.4,
 212 we present results on more practically relevant tasks for scientific discovery, such as the design of
 213 DNA sequences and anti-microbial peptides.

214 4.1 Metrics

215 One core motivation in the conception of GFlowNets, as reported in the original paper [4], was the
216 goal of sampling diverse objects with high-score, according to a reward function.

- 217 • Mean score, as per the highest fidelity oracle f_M , of the top- K samples.
- 218 • Mean pairwise similarity within the top- K samples.

219 Furthermore, since here we are interested in the cost effectiveness of the active learning process, in
220 this section we will evaluate the above metrics as a function of the cost accumulated in querying the
221 multi-fidelity oracles. It is important to note that the multi-fidelity approach is not aimed at achieving
222 *better* mean top- K scores than a single-fidelity active learning counterpart, but rather *the same* mean
223 top- K scores *with a smaller budget*.

224 4.2 Baselines

225 In order to evaluate our approach, and to shed light on the questions stated above, we consider the
226 following baselines:

- 227 • **GFlowNet with highest fidelity (SF-GFN):** GFlowNet based active learning approach from [22]
228 with the highest fidelity oracle to establish a benchmark for performance without considering
229 the cost-accuracy trade-offs.
- 230 • **GFlowNet with random fidelities (Random fid. GFlowNet):** Variant of SF-GFN where the
231 candidates are generated with the GFlowNet but the fidelities are picked randomly and a
232 multi-fidelity acquisition function is used, to investigate the benefit of deciding the fidelity
233 with GFlowNets.
- 234 • **Random candidates and fidelities (Random):** Quasi-random approach where the candidates
235 and fidelities are picked randomly and the top (x, m) pairs scored by the acquisition function
236 are queried.
- 237 • **Multi-fidelity PPO (MF-PPO):** Instantiation of multi-fidelity Bayesian optimization where the
238 acquisition function is optimized using proximal policy optimization [PPO 49].

239 4.3 Synthetic Tasks

240 As an initial assessment of MF-GFNs, we consider two synthetic functions—Branin and Hartmann—
241 widely used in the single- and multi-fidelity Bayesian optimisation literature [44, 53, 28, 36, 16].

242 **Branin** We consider an active learning problem in a two-dimensional space where the target
243 function f_M is the Branin function, as modified in [52] and implemented in `botorch` [2]. We
244 simulate three levels of fidelity, including the true function. The lower-fidelity oracles, the costs of
245 the oracles (0.01, 0.1, 1.0) as well as the number of points queried in the initial training set were
246 adopted from [36]. We provide further details about the task in Appendix B.2. In order to consider a
247 discrete design space, we map the domain to a discrete 100×100 grid. We model this grid with a
248 GFlowNet as in [4, 37]: starting from the origin $(0, 0)$, for any state $s = (x_1, x_2)$, the action space
249 consists of the choice between the exit action or the dimension to increment by 1, provided the next
250 state is in the limits of the grid. Fig. 1a illustrates the results for this task. We observe that MF-GFN
251 is able to reach the minimum of the Branin function with a smaller budget than the single-fidelity
252 counterpart and the baselines.

253 **Hartmann** Next, we consider the 6-dimensional Hartmann function as objective f_M on a hyper-grid
254 domain. As with Branin, we consider three oracles, adopting the lower-fidelity oracles and the set
255 of costs (0.125, 0.25, 1.0) from [53]. We discretize the domain into a six-dimensional hyper-grid of
256 length 10, yielding 10^6 possible candidate points. The results for the task are illustrated in Fig. 1b,
257 which indicate that multi-fidelity active learning with GFlowNets (MF-GFN) offers an advantage
258 over single-fidelity active learning (SF-GFN) as well as some of the other baselines in this higher-
259 dimensional synthetic problem as well. Note that while MF-PPO performs better in this task, as
260 shown in the next experiments, MF-PPO tends to collapse to single modes of the function in more
261 complex high-dimensional scenarios.

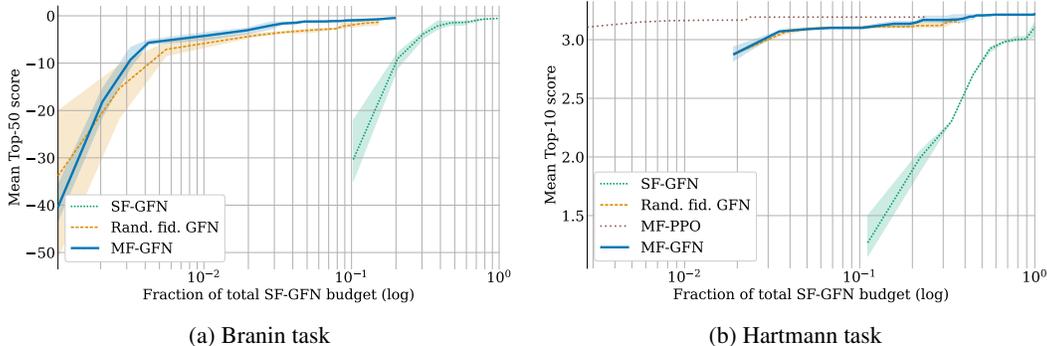


Figure 1: Results on the synthetic tasks—Branin and Hartmann functions. The curves indicate the mean score f_M within the top-50 and top-10 samples (for Branin and Hartmann, respectively) computed at the end of each active learning round and plotted as a function of the budget used. The random baseline is omitted from this plot to facilitate the visualisation since the results were significantly worse in these tasks. We observe that MF-GFN clearly outperforms the single-fidelity counterpart (SF-GFN) and slightly improves upon the GFlowNet baseline that samples random fidelities. On Hartmann, MF-PPO initially outperforms the other methods.

262 4.4 Benchmark Tasks

263 While the synthetic tasks are insightful and convenient for analysis, to obtain a more solid assessment
 264 of the performance of MF-GFN, we evaluate it, together with the other baselines, on more complex,
 265 structured design spaces of practical relevance. We present results on a variety of tasks including DNA
 266 aptamers (Section 4.4.1), anti-microbial peptides (Section 4.4.2) and small molecules (Section 4.4.3).

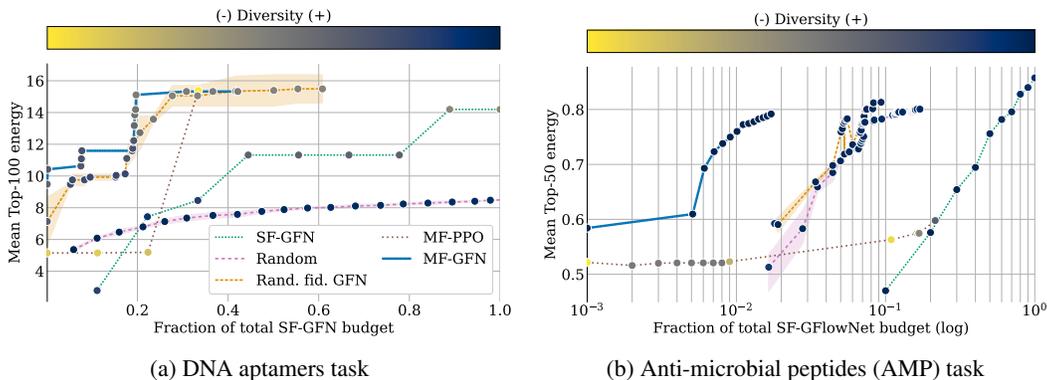


Figure 2: Results on the DNA aptamers and AMP tasks. The curves indicate the mean score f_M within the top-100 and top-50 samples (for DNA and AMP, respectively) computed at the end of each active learning round and plotted as a function of the budget used. The colour of the markers indicates the diversity within the batch (darker colour of the circular dots indicating more diversity). In both the DNA and AMP tasks, MF-GFN outperforms all baselines in terms of cost efficiency, while obtaining great diversity in the final batch of top- K candidates.

267 4.4.1 DNA Aptamers

268 DNA aptamers are single-stranded nucleotide sequences with multiple applications in polymer design
 269 due to their specificity and affinity as sensors in crowded biochemical environments [66, 11, 63, 29].
 270 DNA sequences are represented as strings of nucleobases A, C, T or G. In our experiments, we
 271 consider fixed-length sequences of 30 bases and design a GFlowNet environment where the action
 272 space \mathbb{A} consists of the choice of base to append to the sequence, starting from an empty sequence.
 273 This yields a design space of size $|\mathcal{X}| = 4^{30}$ (ignoring the selection of fidelity in MF-GFN). As
 274 the optimisation objective f_M (highest fidelity) we used the free energy of the secondary structure
 275 as calculated by NUPACK [65]. As a lower fidelity oracle, we trained a transformer model on 1

276 million randomly sampled sequences annotated with f_M , and assigned it a cost $100\times$ smaller than
277 the highest-fidelity oracle. Further details about the task are discussed in Appendix B.4.

278 The main results on the DNA aptamers task are presented in Fig. 2a. We observe that on this task
279 MF-GFN outperforms all other baselines in terms cost efficiency. For instance, MF-GFN achieves
280 the best mean top- K energy achieved by its single-fidelity counterpart with just about 20% of the
281 budget. It is also more efficient than GFlowNet with random fidelities and MF-PPO. Crucially, we
282 also see that MF-GFN maintains a high level of diversity, even after converging to topK reward. On
283 the contrary, MF-PPO is not able to discover diverse samples, as is expected based on prior work [22].

284 4.4.2 Antimicrobial Peptides

285 Antimicrobial peptides are short protein sequences which possess antimicrobial properties. As
286 proteins, these are sequences of amino-acids—a vocabulary of 20 along with a special stop token. We
287 consider variable length proteins sequences with up to 50 residues. We use data from DBAASP [45]
288 containing antimicrobial activity labels, which is split into two sets – one used for training the oracle
289 and one as the initial dataset in the active learning loop, following [22]. To establish the multi-fidelity
290 setting, we train different models with different capacities and with different subsets of the data. The
291 details about these oracles along with additional details about the task are discussed in Appendix B.5.

292 The results in Fig. 2b indicate that even in this task MF-GFN outperforms all other baselines in
293 terms of cost-efficiency. It reaches the same maximum mean top- K score as the random baselines
294 with $10\times$ less budget and almost $100\times$ less budget than SF-GFN. In this task, MF-PPO did not
295 achieve comparable results. Crucially, the diversity of the final batch found by MF-GFN stayed high,
296 satisfying this important criterion in the motivation of this method.

297 4.4.3 Small Molecules

298 Molecules are clouds of interacting electrons (and nuclei) described by a set of quantum mechanical
299 descriptions, or properties. These properties dictate their chemical behaviours and applications.
300 Numerous approximations of these quantum mechanical properties have been developed with different
301 methods at different fidelities, with the famous example of Jacob’s ladder in density functional
302 theory [43]. To demonstrate the capability of MF-GFlowNet to function in the setting of quantum
303 chemistry, we consider MF-GFN out-of-concept tasks in molecular electronic potentials: maximization of
304 adiabatic electron affinity (EA) and (negative) adiabatic ionisation potential (IP). These electronic
305 potentials dictate the molecular redox chemistry, and are key quantities in organic semiconductors,
306 photoredox catalysis, or organometallic synthesis. We employed three oracles that correlate with
307 experimental results as approximations of the scoring function, by uses of varying levels of geometry
308 optimisation to obtain approximations to the adiabatic geometries, followed by the calculation of IP
309 or EA with semi-empirical quantum chemistry XTb (see Appendix) [39]. These three oracles had
310 costs of 1, 3 and 7 (respectively), proportional to their computational running demands. We designed
311 the GFlowNet state space by using sequences of SELFIES tokens [32] (maximum of 64) to represent
312 molecules, starting from an empty sequence; every action consists of appending a new token to the
313 sequence.

314 The realistic configuration and practical relevance of these tasks allow us to draw stronger conclu-
315 sions about the usefulness of multi-fidelity active learning with GFlowNets in scientific discovery
316 applications. As in the other tasks evaluated, we here also found MF-GFN to achieve better cost
317 efficiency at finding high-score top- K molecules (Fig. 3), especially for ionization potentials (Fig. 3a).
318 By clustering the generated molecules, we find that MF-GFN captures as many modes as random
319 generation, far exceeding that of MF-PPO. Indeed, while MF-PPO seems to outperform MF-GFN in
320 the task of electron affinity (Fig. 3b), all generated molecules were from a few clusters, which is of
321 much less utility for chemists.

322 5 Conclusions, Limitations and Future Work

323 In this paper, we present MF-GFN, the first application of GFlowNets for multi-fidelity active learning.
324 Inspired by the encouraging results of GFlowNets in (single-fidelity) active learning for biological
325 sequence design [22] as a method to discover diverse, high-scoring candidates, we propose MF-GFN

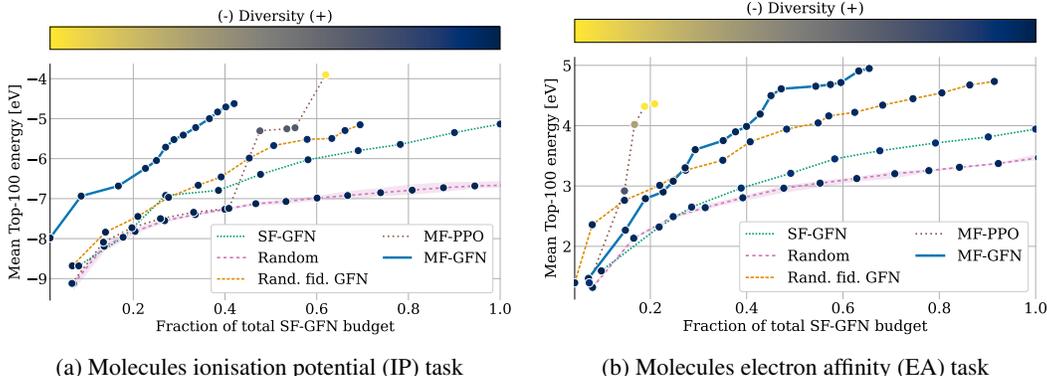


Figure 3: Comparative results on the molecular discovery tasks: (a) ionisation potential (IP), (b) electron affinity (EA). Results illustrate the generally faster convergence of MF-GFN to discover a diverse set of molecules with desirable values of the target property (colour scheme of the circular dots: darker/blue is better than lighter/yellow).

326 to sample the candidates as well as the fidelity at which the candidate is to be evaluated, when
 327 multiple oracles are available with different fidelities and costs.

328 We evaluate the proposed MF-GFN approach in both synthetic tasks commonly used in the multi-
 329 fidelity Bayesian optimization literature and benchmark tasks of practical relevance, such as DNA
 330 aptamer generation, antimicrobial peptide design and molecular modelling. Through comparisons
 331 with previously proposed methods as well as with variants of our method designed to understand the
 332 contributions of different components, we conclude that multi-fidelity active learning with GFlowNets
 333 not only outperforms its single-fidelity active learning counterpart in terms of cost effectiveness and
 334 diversity of sampled candidates, but it also offers an advantage over other multi-fidelity methods due
 335 to its ability to learn a stochastic policy to jointly sample objects and the fidelity of the oracle to be
 336 used to evaluate them.

337 **Broader Impact** Our work is motivated by pressing challenges to sustainability and public health,
 338 and we envision applications of our approach to drug discovery and materials discovery. However,
 339 as with all work on these topics, there is a potential risk of dual use of the technology by nefarious
 340 actors [57].

341 **Limitations and Future Work** Aside from the molecular modelling tasks, our empirical evaluations
 342 in this paper involved simulated oracles with relatively arbitrary costs. Therefore, future work should
 343 evaluate MF-GFN with practical oracles and sets of costs that reflect their computational or financial
 344 demands. Furthermore, we believe a promising avenue that we have not explored in this paper is
 345 the application of MF-GFN in more complex, structured design spaces, such as hybrid (discrete and
 346 continuous) domains [34], as well as multi-fidelity, multi-objective problems [24].

347 References

- 348 [1] Ankit Agrawal and Alok Choudhary. Perspective: Materials informatics and big data: Real-
 349 ization of the “fourth paradigm” of science in materials science. *Apl Materials*, 4(5):053208,
 350 2016.
- 351 [2] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham,
 352 Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo
 353 Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020.
- 354 [3] Ali Bashir, Qin Yang, Jinpeng Wang, Stephan Hoyer, Wenchuan Chou, Cory McLean, Geoff
 355 Davis, Qiang Gong, Zan Armstrong, Junghoon Jang, et al. Machine learning guided aptamer
 356 refinement and discovery. *Nature Communications*, 12(1):2366, 2021.

- 357 [4] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow
358 network based generative models for non-iterative diverse candidate generation. *Advances in*
359 *Neural Information Processing Systems*, 34:27381–27394, 2021.
- 360 [5] Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J. Hu, Mo Tiwari, and Emmanuel Bengio.
361 Gflownet foundations. *arXiv preprint arXiv: Arxiv-2111.09266*, 2021.
- 362 [6] Regine S Bohacek, Colin McMartin, and Wayne C Guida. The art and practice of structure-
363 based drug design: a molecular modeling perspective. *Medicinal research reviews*, 16(1):3–50,
364 1996.
- 365 [7] Keith T Butler, Daniel W Davies, Hugh Cartwright, Olexandr Isayev, and Aron Walsh. Machine
366 learning for molecular and materials science. *Nature*, 559(7715):547–555, 2018.
- 367 [8] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical*
368 *science*, pages 273–304, 1995.
- 369 [9] Konstantinos Chatzilygeroudis, Antoine Cully, Vassilis Vassiliades, and Jean-Baptiste Mouret.
370 Quality-diversity optimization: a novel branch of stochastic optimization. In *Black Box Opti-*
371 *mization, Machine Learning, and No-Free Lunch Theorems*, pages 109–135. Springer, 2021.
- 372 [10] Chi Chen, Yunxing Zuo, Weike Ye, Xiangguo Li, and Shyue Ping Ong. Learning properties
373 of ordered and disordered materials from multi-fidelity data. *Nature Computational Science*,
374 1(1):46–53, 2021.
- 375 [11] David R Corey, Masad J Damha, and Muthiah Manoharan. Challenges and opportunities for
376 nucleic acid therapeutics. *nucleic acid therapeutics*, 32(1):8–13, 2022.
- 377 [12] Payel Das, Tom Sercu, Kahini Wadhawan, Inkit Padhi, Sebastian Gehrmann, Flaviu Cipcigan,
378 Vijil Chenthamarakshan, Hendrik Strobelt, Cicero Dos Santos, Pin-Yu Chen, et al. Accelerated
379 antimicrobial discovery via deep generative models and molecular dynamics simulations. *Nature*
380 *Biomedical Engineering*, 5(6):613–623, 2021.
- 381 [13] Aryan Deshwal and Janardhan Rao Doppa. Combining latent space and structured kernels for
382 bayesian optimization over combinatorial spaces. In *Neural Information Processing Systems*,
383 2021.
- 384 [14] Clyde Fare, Peter Fenner, Matthew Benatan, Alessandro Varsi, and Edward O Pyzer-Knapp. A
385 multi-fidelity machine learning approach to high throughput materials screening. *npj Computa-*
386 *tional Materials*, 8(1):257, 2022.
- 387 [15] Francesco Di Fiore, Michela Nardelli, and Laura Mainini. Active learning and bayesian
388 optimization: a unified perspective to learn with a goal. *arXiv preprint arXiv: Arxiv-2303.01560*,
389 2023.
- 390 [16] Jose Pablo Folch, Robert M Lee, Behrang Shafei, David Walz, Calvin Tsay, Mark van der
391 Wilk, and Ruth Misener. Combining multi-fidelity modelling and asynchronous batch bayesian
392 optimization. *Computers & Chemical Engineering*, 172:108194, 2023.
- 393 [17] Adam Evan Foster. *Variational, Monte Carlo and policy-based approaches to Bayesian*
394 *experimental design*. PhD thesis, University of Oxford, 2021.
- 395 [18] Peter I. Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv: Arxiv-1807.02811*,
396 2018.
- 397 [19] Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.
- 398 [20] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann.
399 Bayesian optimal active search and surveying. *arXiv preprint arXiv:1206.6406*, 2012.
- 400 [21] Antoine Grosnit, Rasul Tutunov, Alexandre Max Maraval, Ryan-Rhys Griffiths, Alexander I.
401 Cowen-Rivers, Lin Yang, Lin Zhu, Wenlong Lyu, Zhitang Chen, Jun Wang, Jan Peters, and
402 Haitham Bou-Ammar. High-dimensional bayesian optimisation with variational autoencoders
403 and deep metric learning. *arXiv preprint arXiv: Arxiv-2106.03609*, 2021.

- 404 [22] Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure FP
405 Dossou, Chanakya Ajit Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, et al.
406 Biological sequence design with gflownets. In *International Conference on Machine Learning*,
407 pages 9786–9801. PMLR, 2022.
- 408 [23] Moksh Jain, Tristan Deleu, Jason Hartford, Cheng-Hao Liu, Alex Hernandez-Garcia, and
409 Yoshua Bengio. Gflownets for ai-driven scientific discovery. *Digital Discovery*, 2023.
- 410 [24] Moksh Jain, Sharath Chandra Raparthy, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Yoshua
411 Bengio, Santiago Miret, and Emmanuel Bengio. Multi-objective gflownets. *arXiv preprint*
412 *arXiv:2210.12765*, 2022.
- 413 [25] Shali Jiang, Roman Garnett, and Benjamin Moseley. Cost effective active search. *Advances in*
414 *Neural Information Processing Systems*, 32, 2019.
- 415 [26] Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, and
416 Roman Garnett. Efficient nonmyopic active search. In *International Conference on Machine*
417 *Learning*, pages 1714–1723. PMLR, 2017.
- 418 [27] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ron-
419 neberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al.
420 Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- 421 [28] Kirthevasan Kandasamy, Gautam Dasarathy, Junier B. Oliva, Jeff Schneider, and Barnabas
422 Poczos. Multi-fidelity gaussian process bandit optimisation, 2019.
- 423 [29] Michael Kilgour, Tao Liu, Brandon D Walker, Pengyu Ren, and Lena Simine. E2edna: Simula-
424 tion protocol for dna aptamers with ligands. *Journal of Chemical Information and Modeling*,
425 61(9):4139–4144, 2021.
- 426 [30] Ross D King, Kenneth E Whelan, Ffion M Jones, Philip GK Reiser, Christopher H Bryant,
427 Stephen H Muggleton, Douglas B Kell, and Stephen G Oliver. Functional genomic hypothesis
428 generation and experimentation by a robot scientist. *Nature*, 427(6971):247–252, 2004.
- 429 [31] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- 430 [32] Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik.
431 Self-referencing embedded strings (selfies): A 100% robust molecular string representation.
432 *Machine Learning: Science and Technology*, 1(4):045024, 2020.
- 433 [33] A Gilad Kusne, Heshan Yu, Changming Wu, Huairuo Zhang, Jason Hattrick-Simpers, Brian
434 DeCost, Suchismita Sarker, Corey Oses, Cormac Toher, Stefano Curtarolo, et al. On-the-
435 fly closed-loop materials discovery via bayesian active learning. *Nature communications*,
436 11(1):5966, 2020.
- 437 [34] Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghuai Zhang, Alexandra Volokhova, Alex
438 Hernández-García, Léna Néhale Ezzine, Yoshua Bengio, and Nikolay Malkin. A theory of
439 continuous generative flow networks. *International Conference on Machine Learning*, 2023.
- 440 [35] Shibo Li, Robert M Kirby, and Shandian Zhe. Deep multi-fidelity active learning of high-
441 dimensional outputs. *arXiv preprint arXiv:2012.00901*, 2020.
- 442 [36] Shibo Li, Wei Xing, Mike Kirby, and Shandian Zhe. Multi-fidelity bayesian optimization via
443 deep neural networks. *Advances in Neural Information Processing Systems*, 2020.
- 444 [37] Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory
445 balance: Improved credit assignment in gflownets, 2022.
- 446 [38] Henry B. Moss, David S. Leslie, Javier Gonzalez, and Paul Rayson. Gibbon: General-purpose
447 information-based bayesian optimisation, 2021.
- 448 [39] Hagen Neugebauer, Fabian Bohle, Markus Bursch, Andreas Hansen, and Stefan Grimme.
449 Benchmark study of electrochemical redox potentials calculated with semiempirical and dft
450 methods. *The Journal of Physical Chemistry A*, 124(35):7166–7176, 2020.

- 451 [40] Quan Nguyen, Arghavan Modiri, and Roman Garnett. Nonmyopic multifidelity active search.
452 In *International Conference on Machine Learning*, pages 8109–8118. PMLR, 2021.
- 453 [41] Robert G Parr. Density functional theory of atoms and molecules. In *Horizons of Quantum*
454 *Chemistry: Proceedings of the Third International Congress of Quantum Chemistry Held at*
455 *Kyoto, Japan, October 29–November 3, 1979*, pages 5–15. Springer, 1980.
- 456 [42] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods
457 in uncertainty propagation, inference, and optimization. *Siam Review*, 60(3):550–591, 2018.
- 458 [43] John P Perdew and Karla Schmidt. Jacob’s ladder of density functional approximations for
459 the exchange-correlation energy. In *AIP Conference Proceedings*, volume 577, pages 1–20.
460 American Institute of Physics, 2001.
- 461 [44] Paris Perdikaris, M. Raissi, Andreas C. Damianou, ND Lawrence, and George Em Karniadakis.
462 Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings*
463 *of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473, 2017.
- 464 [45] Malak Pirtskhalava, Anthony A Armstrong, Maia Grigolava, Mindia Chubinidze, Evgenia Alim-
465 barashvili, Boris Vishnepolsky, Andrei Gabrielian, Alex Rosenthal, Darrell E Hurt, and Michael
466 Tartakovsky. Dbasp v3: database of antimicrobial/cytotoxic activity and structure of peptides
467 as a resource for development of new therapeutics. *Nucleic acids research*, 49(D1):D288–D297,
468 2021.
- 469 [46] Pavel G Polishchuk, Timur I Madzhidov, and Alexandre Varnek. Estimation of the size of
470 drug-like chemical space based on gdb-17 data. *Journal of computer-aided molecular design*,
471 27:675–679, 2013.
- 472 [47] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine*
473 *Learning*. The MIT Press, 11 2005.
- 474 [48] Herbert E. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the*
475 *American Mathematical Society*, 58:527–535, 1952.
- 476 [49] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal
477 policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 478 [50] Burr Settles. Active learning literature survey. *Independent Technical Report*, 2009.
- 479 [51] David S Sholl and Janice A Steckel. *Density functional theory: a practical introduction*. John
480 Wiley & Sons, 2022.
- 481 [52] András Sobester, Alexander Forrester, and Andy Keane. *Appendix: Example Problems*, pages
482 195–203. John Wiley & Sons, Ltd, 2008.
- 483 [53] Jialin Song, Yuxin Chen, and Yisong Yue. A general framework for multi-fidelity bayesian
484 optimization with gaussian processes. In *International Conference on Artificial Intelligence*
485 *and Statistics*, 2018.
- 486 [54] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process
487 optimization in the bandit setting: No regret and experimental design. *arXiv preprint*
488 *arXiv:0912.3995*, 2009.
- 489 [55] Jonathan M Stokes, Kevin Yang, Kyle Swanson, Wengong Jin, Andres Cubillos-Ruiz, Nina M
490 Donghia, Craig R MacNair, Shawn French, Lindsey A Carfrae, Zohar Bloom-Ackermann, et al.
491 A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702, 2020.
- 492 [56] Shion Takeno, Hitoshi Fukuoka, Yuhki Tsukada, Toshiyuki Koyama, Motoki Shiga, Ichiro
493 Takeuchi, and Masayuki Karasuyama. Multi-fidelity Bayesian optimization with max-value
494 entropy search and its parallelization. In Hal Daumé III and Aarti Singh, editors, *Proceedings*
495 *of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of*
496 *Machine Learning Research*, pages 9334–9345. PMLR, 13–18 Jul 2020.
- 497 [57] Fabio Urbina, Filippa Lentzos, Cédric Invernizzi, and Sean Ekins. Dual use of artificial-
498 intelligence-powered drug discovery. *Nature Machine Intelligence*, 4(3):189–191, 2022.

- 499 [58] Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization,
500 2018.
- 501 [59] Manfred KK Warmuth, Gunnar Rätsch, Michael Mathieson, Jun Liao, and Christian Lemmen.
502 Active learning in the drug discovery process. *Advances in Neural information processing*
503 *systems*, 14, 2001.
- 504 [60] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel
505 learning. In *Artificial intelligence and statistics*, pages 370–378. PMLR, 2016.
- 506 [61] Jian Wu, Saul Toscano-Palmerin, Peter I. Frazier, and Andrew Gordon Wilson. Practical
507 multi-fidelity bayesian optimization for hyperparameter tuning, 2019.
- 508 [62] Dezhen Xue, Prasanna V Balachandran, John Hogden, James Theiler, Deqing Xue, and Turab
509 Lookman. Accelerated search for materials with targeted properties by adaptive design. *Nature*
510 *communications*, 7(1):1–9, 2016.
- 511 [63] Joseph D Yesselman, Daniel Eiler, Erik D Carlson, Michael R Gotrik, Anne E d’Aquino,
512 Alexandra N Ooms, Wipapat Kladwang, Paul D Carlson, Xuesong Shi, David A Costantino, et al.
513 Computational design of three-dimensional rna structure and function. *Nature nanotechnology*,
514 14(9):866–873, 2019.
- 515 [64] Ruihao Yuan, Zhen Liu, Prasanna V Balachandran, Deqing Xue, Yumei Zhou, Xiangdong Ding,
516 Jun Sun, Dezhen Xue, and Turab Lookman. Accelerated discovery of large electrostrains in
517 batio₃-based piezoelectrics using active learning. *Advanced materials*, 30(7):1702884, 2018.
- 518 [65] Joseph N Zadeh, Conrad D Steenberg, Justin S Bois, Brian R Wolfe, Marshall B Pierce, Asif R
519 Khan, Robert M Dirks, and Niles A Pierce. Nupack: Analysis and design of nucleic acid
520 systems. *Journal of computational chemistry*, 32(1):170–173, 2011.
- 521 [66] Wenhui Zhou, Runjhun Saran, and Juewen Liu. Metal sensing by dna. *Chemical reviews*,
522 117(12):8272–8325, 2017.
- 523 [67] C Lawrence Zitnick, Lowik Chanussot, Abhishek Das, Siddharth Goyal, Javier Heras-Domingo,
524 Caleb Ho, Weihua Hu, Thibaut Lavril, Aini Palizhati, Morgane Riviere, et al. An introduction
525 to electrocatalyst design using machine learning for renewable energy storage. *arXiv preprint*
526 *arXiv:2010.09435*, 2020.