Amortized Bayesian Experimental Design for Decision-Making

Daolang Huang

Aalto University daolang.huang@aalto.fi

Luigi Acerbi

University of Helsinki luigi.acerbi@helsinki.fi

Yujia Guo

Aalto University yujia.guo@aalto.fi

Samuel Kaski

Aalto University University of Manchester samuel.kaski@aalto.fi

Abstract

Many critical decisions, such as personalized medical diagnoses and product pricing, are made based on insights gained from designing, observing, and analyzing a series of experiments. This highlights the crucial role of experimental design, which goes beyond merely collecting information on system parameters as in traditional Bayesian experimental design (BED), but also plays a key part in facilitating downstream decision-making. Most recent BED methods use an amortized policy network to rapidly design experiments. However, the information gathered through these methods is suboptimal for down-the-line decision-making, as the experiments are not inherently designed with downstream objectives in mind. In this paper, we present an amortized decision-aware BED framework that prioritizes maximizing downstream decision utility. We introduce a novel architecture, the Transformer Neural Decision Process (TNDP), capable of instantly proposing the next experimental design, whilst inferring the downstream decision, thus effectively amortizing both tasks within a unified workflow. We demonstrate the performance of our method across several tasks, showing that it can deliver informative designs and facilitate accurate decision-making.

1 Introduction

In a wide array of disciplines, from clinical trials (Cheng and Shen, 2005) to medical imaging (Burger et al., 2021), a fundamental challenge is the design of experiments to infer the distribution of some unobservable, unknown properties of the systems under study. *Bayesian Experimental Design* (BED) (Lindley, 1956; Chaloner and Verdinelli, 1995; Ryan et al., 2016; Rainforth et al., 2024) is a powerful framework in this context, guiding and optimizing experimental design by maximizing the expected amount of information about parameters gained from experiments, see Fig. 1(a). However, the ultimate goal in many tasks extends beyond parameter inference to inform a *downstream decision-making* task by leveraging our understanding of these parameters. For example, in personalized medical diagnostics, a model is built based on historical data to facilitate an optimal treatment for a new patient (Bica et al., 2021). This data typically comprises patient covariates, administered treatments, and observed outcomes. Since the query of such data tends to be expensive due to, e.g., privacy issues, we need to actively design queries to optimize resource utilization. Here, when the goal is to improve decisions, the strategy of experimental designs should not focus solely on inferring the parameters of the model, but rather on guiding the final decision-making for the new patient, to ensure that each query contributes maximally to the diagnostic decision.

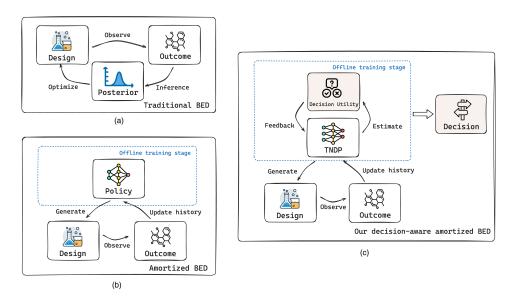


Figure 1: **Overview of BED workflows.** (a) Traditional BED, which iterates between optimizing designs, running experiments, and updating the model via Bayesian inference. (b) Amortized BED, which uses a policy network for rapid experimental design generation. (c) Our decision-aware amortized BED integrates decision utility in training to facilitate downstream decision-making.

Traditional BED methods (Rainforth et al., 2018; Foster et al., 2019, 2020; Kleinegesse and Gutmann, 2020) do not take down-the-line decision-making tasks into account during the experimental design phase. As a result, inference and decision-making processes are carried out separately, which is suboptimal for decision-making in scenarios where experiments can be adaptively designed. *Loss-calibrated inference*, which was originally introduced by Lacoste-Julien et al. (Lacoste-Julien et al., 2011) for variational approximations in Bayesian inference, provides a concept that adjusts the inference process to capture posterior regions critical for decision-making tasks. Rather than focusing on parameter estimation, the idea is to directly maximize the expected downstream utility, recognizing that accurate decision-making can proceed without exact knowledge of the posterior as not all regions of the posterior contribute equally to the downstream task. Inspired by this concept, we could consider integrating decision-making directly into the experimental design process to align the proposed designs more closely with the ultimate decision-making task.

To pick the next optimal design, standard BED methods require estimating and optimizing the expected information gain (EIG) over the design space, which can be extremely time-consuming. This limitation has led to the development of *amortized* BED (Foster et al., 2021; Ivanova et al., 2021; Blau et al., 2022, 2023), a *policy-based* method which leverages a neural network policy trained on simulated experimental trajectories to quickly generate designs, as illustrated in Fig. 1(b). Given an experiment history, these policy-based methods determine the next experimental design through a single forward pass, significantly speeding up the design process. Another significant advantage of amortized BED methods is their capacity to extract and utilize unstructured domain knowledge embedded in historical data. Unlike traditional methods, which never reuse the information from past data, amortized methods can integrate this knowledge to refine and improve design strategies for new experiments. In our problem setting, the benefits of amortization are also valuable where decisions must be made swiftly, such as when determining optimal treatment for patients in urgent settings.

In this paper, we propose an amortized decision-making-aware BED framework, see Fig. 1(c). We identify two key aspects where previous amortized BED methods fall short when applied to downstream decision-making tasks. First, the training objective of the existing methods does not consider downstream decision tasks. Therefore, we introduce the concept of *Decision Utility Gain* (DUG) to guide experimental design to better align with the downstream objective. DUG is designed to measure the improvement in the *maximum* expected utility derived from the new experiment. Second, to obtain the optimal decision, we still need to explicitly approximate the predictive distribution of the outcomes to estimate the utility. Current amortized methods learn this distribution only implicitly and therefore require extra computation for the decision-making process. To address this, we propose

a novel *Transformer neural decision process* (TNDP) architecture with dual output heads: one acting as a policy network to propose new designs, and another to approximate the predictive distribution to support the downstream decision-making. This setup allows an iterative approach where the system autonomously proposes informative experimental designs and makes optimal final decisions. Finally, since our ultimate goal is to make optimal decisions at the final stage, which may involve multiple experiments, it is crucial that our experimental designs are not myopic or overly greedy by only maximizing next-step decision utility. Therefore, we develop a non-myopic objective function that ensures decisions are made with a comprehensive consideration of future outcomes.

Contributions. In summary, the contributions of our work include:

- We propose the concept of *decision utility gain* (DUG) for guiding the next experimental design with a direct focus on optimizing down-the-line decision-making tasks.
- We present a novel architecture *Transformer neural decision process* (TNDP) designed to amortize the experimental designs directly for downstream decision-making.
- We empirically show the effectiveness of TNDP across a variety of experimental design tasks involving decision-making, where it significantly outperforms other methods that do not consider downstream decisions.

2 Preliminaries

2.1 Bayesian experimental design

BED (Lindley, 1956; Ryan et al., 2016; Rainforth et al., 2024) is a powerful statistical framework that optimizes the experimental design process. The primary goal of BED is to sequentially select a set of experimental designs $\xi \in \Xi$ and gather outcomes y, to maximize the amount of information obtained about the parameters of interest, denoted as θ . Essentially, BED seeks to minimize the entropy of the posterior distribution of θ or, equivalently, to maximize the information that the experimental outcomes provide about θ .

At the heart of BED lies the concept of *Expected Information Gain* (EIG), which quantifies the value of different experimental designs based on how much they are expected to reduce uncertainty about θ , measured in terms of expected entropy $(H[\cdot])$ reduction:

$$EIG(\xi) = \mathbb{E}_{p(y|\xi)} \left[H \left[p(\theta) \right] - H \left[p(\theta|\xi, y) \right] \right]. \tag{1}$$

The optimal design is then defined as $\xi^* = \arg\max_{\xi \in \Xi} \mathrm{EIG}(\xi)$. In practice, calculating EIG is a computationally challenging task because it involves integrations over $p(y|\xi)$ and $p(\theta|\xi,y)$, which are both intractable. In recent years, various methods have been proposed to make the computation of the EIG feasible in practical scenarios, such as nested Monte Carlo estimators (Rainforth et al., 2018) and variational approximations (Foster et al., 2019). However, even with these advancements, the computational load remains significant, hindering feasibility in tasks that demand rapid designs. This limitation has pushed forward the development of amortized BED methods, which significantly reduce computational demands during the deployment stage.

2.2 Amortized BED

Amortized BED methods represent a significant shift from traditional experimental optimization techniques. Instead of optimizing for each experimental design separately, Foster et al. (2021) developed a parameterized policy π , which directly maps the experimental history $h_{1:t} = \{(\xi_1, y_1), ..., (\xi_t, y_t)\}$ to the next design $\xi_{t+1} = \pi(h_{1:t})$. To train such a policy network, Foster et al. (2021) proposed using sequential Prior Contrastive Estimation (sPCE) to optimize the lower bound of the total EIG across the entire T-step experiments trajectory:

$$sPCE(\pi, L) = \mathbb{E}_{p(\theta_{0:L})p(h_{1:T}|\theta_0, \pi)} \left[\log \frac{p(h_{1:T}|\theta_0, \pi)}{\frac{1}{L+1} \sum_{\ell=0}^{L} p(h_{1:T}|\theta_\ell, \pi)} \right], \tag{2}$$

where $\theta_{1:L}$ are contrastive samples drawn from the prior distribution. Although the training of such a policy network is computationally expensive, once trained, the network can act as an oracle to quickly propose the next design through a single forward pass, thus amortizing the initial training cost over numerous deployments.

2.3 Bayesian decision theory

Bayesian decision theory (Berger, 2013) provides an axiomatic framework for decision-making under uncertainty, systematically incorporating probabilistic beliefs about unknown parameters into decision-making processes. It introduces a task-specific utility function $u(\theta, a)$, which quantifies the value of the outcomes from different decisions $a \in \mathcal{A}$ when the system is in state θ . The optimal decision is then determined by maximizing the expected utility, which integrates the utility function over the unknown system parameters, given the available knowledge $h_{1:t}$:

$$a^* = \operatorname*{max}_{a \in A} \mathbb{E}_{p(\theta|h_{1:t})}[u(\theta, a)]. \tag{3}$$

In many scenarios, outcomes are stochastic and it is more typical to make decisions based on their predictive distribution $p(y|\xi,h_{1:t}) = \mathbb{E}_{p(\theta|h_{1:t})}[p(y|\xi,\theta,h_{1:t})]$, such as in clinical trials where the optimal treatment is chosen based on predicted patient responses rather than solely on underlying biological mechanisms. A similar setup can be found in (Kuśmierczyk et al., 2019; Vadera et al., 2021). As we switch the belief about the state of the system to the outcomes and to keep as much information as possible, we need to evaluate the effect of θ on all points of the design space. Thus, instead of the posterior over the latent state θ , we represent our belief directly as $p(y_{\Xi}|h_{1:t}) \equiv \{p(y|\xi,h_{1:t})\}_{\xi \in \Xi}$, i.e. a joint predictive (posterior) distribution of outcomes over all possible designs given the current information $h_{1:t}$. The utility is then expressed as $u(y_{\Xi},a)$, which relies on the decision a and all possible predicted outcomes y_{Ξ} . It is a natural extension of the traditional definition of utility by marginalizing out the posterior distribution of θ . The rule of making the optimal decision is reformulated in terms of the predictive distribution as:

$$a^* = \arg\max_{a \in A} \mathbb{E}_{p(y_{\Xi}|h_{1:t})}[u(y_{\Xi}, a)]. \tag{4}$$

Traditional methods usually separate the inference and decision-making steps, which are optimal when the true posterior or the predictive distribution can be computed exactly. However, in most cases the posteriors are not directly accessible, and we often resort to using approximate distributions. This necessity results in a suboptimal decision-making process as the approximate posteriors often focus on representing the full posterior yet fail to ensure high accuracy in regions crucial for decision-making. *Loss-calibrated inference* (Lacoste-Julien et al., 2011) emerges as an effective solution to address this problem. It calibrates the inference by focusing on utility rather than mere accuracy of the approximation, thereby ensuring a more targeted posterior estimation. This method has been applied to improving Markov chain Monte Carlo (MCMC) methods (Abbasnejad et al., 2015), Bayesian neural networks (Cobb et al., 2018) and expectation propagation (Morais and Pillow, 2022).

3 Decision-aware BED

3.1 Problem setup

Our objective is to optimize the experimental design process for down-the-line decision-making. In this paper, we consider scenarios in which we design a series of experiments $\xi \in \Xi$ and observe corresponding outcomes y to inform a final decision-making step. We assume we have a fixed experimental budget with T query steps. For decision-making, we consider a set of possible decisions, denoted as $a \in \mathcal{A}$, with the objective of identifying an optimal decision a^* that maximizes a predefined prediction-based utility function $u(y_\Xi, a)$.

3.2 Decision Utility Gain

Our method focuses on designing the experiments to directly improve the quality of the final decision-making. To quantify the effectiveness of each experimental design in terms of decision-making, we introduce *Decision Utility Gain* (DUG), which is defined as the difference in the expected utility of the best decision, with the new information obtained from the current experimental design, versus the best decision with the information obtained from previous experiments.

¹This definition assumes conditional independence of the outcomes given the design. More generally, $p(y_{\Xi}|h_{1:t})$ defines a joint distribution or a *stochastic process* indexed by the set Ξ (Parzen, 1999), where a familiar example could be a Gaussian process posterior defined on $\Xi \subseteq \mathbb{R}^d$ (Rasmussen and Williams, 2006).

Definition 3.1. Given a historical experimental trajectory $h_{1:t-1}$, the *Decision Utility Gain* (DUG) for a given design ξ_t and its corresponding outcome y_t at step t is defined as follows:

$$DUG(\xi_t, y_t) = \max_{a \in A} \mathbb{E}_{p(y_\Xi | h_{1:t-1} \cup \{(\xi_t, y_t)\})} [u(y_\Xi, a)] - \max_{a \in A} \mathbb{E}_{p(y_\Xi | h_{1:t-1})} [u(y_\Xi, a)].$$
 (5)

DUG measures the improvement in the *maximum* expected utility from observing a new experimental design, differing in this from standard marginal utility gain (see e.g., Garnett, 2023). The optimal design is the one that provides the largest increase in maximal expected utility. Shifting from parameter-centric to utility-centric evaluation, we directly evaluate the design's influence on the decision utility, bypassing the need to reduce the uncertainty of unknown latent parameters.

At the time we choose the design ξ_t , the outcome remains uncertain. Therefore, we should consider the *Expected Decision Utility Gain* (EDUG) with respect to the marginal distribution of the outcomes to select the next design.

Definition 3.2. The *Expected Decision Utility Gain* (EDUG) for a design ξ_t , given the historical experimental trajectory $h_{1:t-1}$, is defined as:

$$EDUG(\xi_t) = \mathbb{E}_{p(y_t|\xi_t, h_{1:t-1})} \left[DUG(\xi_t, y_t) \right]. \tag{6}$$

With EDUG, we can guide the experimental design without calculating the posterior distribution. If we knew the true predictive distribution, we could always determine the one-step lookahead optimal design by maximizing EDUG across the design space with $\xi^* = \arg\max_{\xi \in \Xi} \mathrm{EDUG}(\xi)$. However, in practice, the true predictive distributions are often unknown, making the optimization of EDUG exceptionally challenging. This difficulty arises due to the inherent bi-level optimization problem and the need to evaluate two layers of expectations, both over the unknown predictive distribution.

To avoid the expensive computational cost of optimizing EDUG, we propose leveraging a policy network, inspired by Foster et al. (2021), that directly maps historical data to the next design. This approach sidesteps the continuous need to optimize EDUG by learning a design strategy over many simulated experiment trajectories during the training phase. It can dramatically reduce computational demands at deployment, allowing for more efficient real-time decisions.

4 Amortizing decision-aware BED

A fully amortized BED framework for decision-making requires not only amortizing the experimental design but also the predictive distribution to approximate the expected utility. Moreover, *permutation invariance* is often assumed in sequential BED (Foster et al., 2021)², meaning that the sequence of experiments does not influence the cumulative information gained. Conditional neural processes (CNPs) (Garnelo et al., 2018; Nguyen and Grover, 2022; Huang et al., 2023b) provide a suitable basis for developing our framework due to their design, which not only respects the permutation invariance of the inputs by treating them as an unordered set but also amortizes modeling of the predictive distributions. See Appendix A for a brief introduction to CNPs and TNPs.

4.1 Transformer Neural Decision Processes

The architecture of our model, termed *Transformer Neural Decision Process* (TNDP), is a novel architecture building upon the Transformer neural process (TNP) (Nguyen and Grover, 2022). It aims to amortize both the experimental design and the predictive distributions for the subsequent decision-making process. The data architecture of our system comprises four parts $D = \{D^{(c)}, D^{(p)}, D^{(q)}, GI\}$:

- A context set $D^{(c)} = h_{1:t} = \{(\xi_i^{(c)}, y_i^{(c)})\}_{i=1}^t$ contains all past t-step designs and outcomes.
- A **prediction set** $D^{(p)} = \{(\xi_i^{(p)}, y_i^{(p)})\}_{i=1}^{n_p}$ consists of n_p design-outcome pairs used for approximating $p(y_{\Xi}|h_{1:t})$, which is closely related to the training objective of the CNPs. The output from this head can then be used to estimate the expected utility.
- A query set $D^{(q)} = \{\xi_i^{(q)}\}_{i=1}^{n_q}$ consists of n_q candidate experimental designs being considered for the next step. In scenarios where the design space Ξ is continuous, we randomly sample a set of query points for each iteration during training. In the deployment phase, optimal experimental designs can be obtained by optimizing the model's output.

²When permutation invariance does not hold in some cases, our model can be easily adapted by adding positional encoding to the input.

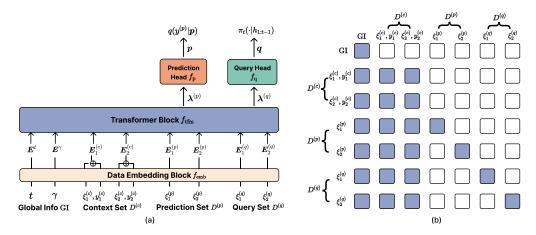


Figure 2: **Illustration of TNDP.** (a) An overview of TNDP architecture with input consisting of 2 observed design-outcome pairs from $D^{(c)}$, 2 designs from $D^{(p)}$ for prediction, and 2 candidate designs from $D^{(q)}$ for query. (b) The corresponding attention mask. The colored squares indicate that the elements on the left can attend to the elements on the top in the self-attention layer of f_{tfm} .

• Global information GI = $[t, \gamma]$ where t represents the current step in the experimental sequence, and γ encapsulates task-related information, which could include contextual data relevant to the decision-making process. We will further explain the choice of γ in Section 6.

TNDP comprises four main components: the data embedder block $f_{\rm emb}$, the Transformer block $f_{\rm tfm}$, the prediction head $f_{\rm p}$, and the query head $f_{\rm q}$. Each component plays a distinct role in the overall decision-aware BED process. The full architecture is shown in Fig. 2(a).

At first, each set of D is processed by the **data embedder block** $f_{\rm emb}$ to map to an aligned embedding space. These embeddings are then concatenated to form a unified representation $E = {\rm concat}(E^{(c)}, E^{(p)}, E^{(q)}, E^{\rm GI})$. Please refer to Appendix B for a detailed explanation of how we embed the data. After the initial embedding, the **Transformer block** $f_{\rm tfm}$ processes E using self-attention mechanisms to produce a single attention matrix, which is subsequently processed by an attention mask (see Fig. 2(b)) that allows for selective interactions between different data components, ensuring that each part contributes appropriately to the final output. To explain, each design from the prediction set $D^{(p)}$ is configured to attend to itself, the global information, and the historical data, reflecting the dependence of the predictions on the historical data and the independence from other designs. Similarly, each $\xi^{(q)}$ in the query set $D^{(q)}$ is also restricted to attend only to itself, the global information, and the historical data. This setup preserves the independence of each candidate design, ensuring that the evaluation of one design neither influences nor is influenced by others. The output of $f_{\rm tfm}$ is then split according to the specific needs of the query and prediction head $\lambda = [\lambda^{(p)}, \lambda^{(q)}] = f_{\rm tfm}(E)$.

The primary role of the **prediction head** f_p is to approximate $p(y_{\Xi}|h_{1:t})$ with a family of parameterized distributions $q(y_{\Xi}|p_t)$, where $p_t = f_p(\lambda_t^{(p)})$ is the output of f_p at the step t. The training of f_p is by minimizing the negative log-likelihood of the predicted probabilities:

$$\mathcal{L}^{(p)} = -\sum_{t=1}^{T} \sum_{i=1}^{n_p} \log q(y_i^{(p)} | \boldsymbol{p}_{i,t}) = -\sum_{t=1}^{T} \sum_{i=1}^{n_p} \log \mathcal{N}(y_i^{(p)} | \boldsymbol{\mu}_{i,t}, \boldsymbol{\sigma}_{i,t}^2), \tag{7}$$

where $p_{i,t}$ represents the output of design $\xi_i^{(p)}$ at step t. Here we choose a Gaussian likelihood with μ and σ representing the predicted mean and standard deviation split from p.

The query head f_q processes the transformed embeddings $\lambda^{(q)}$ from the Transformer block to generate a policy distribution over possible experimental designs. Specifically, it converts the embeddings into a probability distribution used to select the next experimental design. The outputs of the query head, $q = f_q(\lambda^{(q)})$, are mapped to a probability distribution via a Softmax function:

$$\pi(\xi_{j,t}^{(q)}|h_{1:t-1}) = \frac{\exp(\mathbf{q}_{j,t})}{\sum_{i=0}^{n_q} \exp(\mathbf{q}_{i,t})},$$
(8)

where $q_{j,t}$ represents the t-step's output for the candidate design $\xi_j^{(q)}$.

To design a reward signal that guides f_q in proposing informative designs, we first define a single-step immediate reward based on DUG (Eq. (5)), replacing the true predictive distribution with our approximated distribution:

$$r_t(\xi_t^{(q)}) = \max_{a \in A} \mathbb{E}_{q(y_\Xi|\mathbf{p}_t)} \left[u(y_\Xi, a) \right] - \max_{a \in A} \mathbb{E}_{q(y_\Xi|\mathbf{p}_{t-1})} \left[u(y_\Xi, a) \right]. \tag{9}$$

This reward quantifies how the experimental design influences our decision-making by estimating the improvement in expected utility that results from incorporating new experimental outcomes. However, this objective remains myopic, as it does not account for the future or the final decision-making. To address this, we employ the REINFORCE algorithm (Williams, 1992), which allows us to consider the impact of the current design on future rewards. The final loss of $f_{\rm q}$ can be written as the negative expected reward for the complete experimental trajectory:

$$\mathcal{L}^{(q)} = -\sum_{t=1}^{T} R_t \log \pi(\xi_t^{(q)} | h_{1:t-1}), \tag{10}$$

where $R_t = \sum_{k=t}^T \alpha^{k-t} r_k(\xi_k^{(q)})$ represents the non-myopic discounted reward. The discount factor α is used to decrease the importance of rewards received at later time step. $\xi_t^{(q)}$ is obtained through sampling from the policy distribution $\xi_t^{(q)} \sim \pi(\cdot|h_{1:t-1})$.

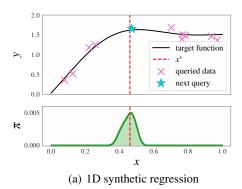
The update of $f_{\rm q}$ depends critically on the accuracy with which $f_{\rm p}$ approximates the predictive distribution. Ultimately, the effectiveness of decision-making relies on the informativeness of the designs proposed by $f_{\rm q}$, ensuring that every step in the experimental trajectory is optimally aligned with the overarching goal of maximizing the decision utility. The full algorithm of our method is shown in Appendix C.

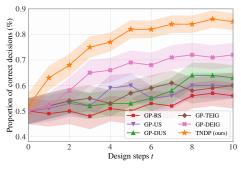
5 Related work

Lindley (1972) proposes the first decision-theoretic BED framework, later reiterated by Chaloner and Verdinelli (1995). However, their utility is defined based on individual designs, while our utility is formulated in terms of a stochastic process and is designed for the final decision-making task after multiple rounds of experimental design. Recently, several other BED frameworks that focus on different downstream properties have been proposed. Bayesian Algorithm Execution (BAX) (Neiswanger et al., 2021) develops a BED framework that optimizes experiments based on downstream properties of interest. BAX introduces a new metric that queries the next experiment by maximizing the mutual information between the property and the outcome. CO-BED (Ivanova et al., 2023) introduces a contextual optimization method within the BED framework, where the design phase incorporates information-theoretic objectives specifically targeted at optimizing contextual rewards. Neiswanger et al. (2022) presents an information-based acquisition function for Bayesian optimization which explicitly considers the downstream task. Zhong et al. (2024) proposes a goal-oriented BED framework for nonlinear models using MCMC to optimize the EIG on predictive quantities of interest. Filstroff et al. (2024) presents a framework for active learning that queries data to reduce the uncertainty on the posterior distribution of the optimal downstream decision.

In recent years, various amortized BED methods have emerged. Foster et al. (2021) is the first to introduce this framework; subsequent work extends it to scenarios with unknown likelihood (Ivanova et al., 2021) and improved performance using reinforcement learning (Blau et al., 2022; Lim et al., 2022). The latest research proposes a semi-amortized framework that periodically updates the policy during the experiment to improve adaptability (Ivanova et al., 2024). Maraval et al. (2024) proposes a fully amortized Bayesian optimization (BO) framework that employs a similar TNP architecture, while their work focuses specifically on BO objectives, our approach addresses general downstream decision-making tasks. Additionally, our framework introduces a novel coupled training objective between query and prediction heads, providing a more integrated architecture for decision-making.

Our proposed architecture is based on pre-trained Transformer models. Transformer-based neural processes (Müller et al., 2021; Nguyen and Grover, 2022; Chang et al., 2024) serve as the foundational structure for our approach, but they have not considered experimental design. Decision Transformers (Chen et al., 2021; Zheng et al., 2022) can be used for sequentially designing experiments. However, we additionally amortize the predictive distribution, making the learning process more challenging.





(b) Decision-aware active learning

Figure 3: **Results of synthetic regression and decision-aware active learning.** (a) The top figure represents the true function and the initial known points. The red line indicates the location of x^* . The blue star marks the next query point, sampled from the policy's predicted distribution shown in the bottom figure. (b) Mean and standard error of the proportion of correct decisions on 100 test points w.r.t. the acquisition steps. Our TNDP significantly outperforms other methods.

6 Experiments

In this section, we evaluate our proposed framework on several tasks. Our experimental approach is detailed in Appendix B. In Appendix F.3, we provide additional ablation studies of TNDP to show the effectiveness of our query head and the non-myopic objective function. The code to reproduce our experiments is available at https://github.com/huangdaolang/amortized-decision-aware-bed.

6.1 Toy example: synthetic regression

We begin with an illustrative example to show how our TNDP works. We consider a 1D synthetic regression task where the goal is to perform regression at a specific test point x^* on an unknown function. To accurately predict this point, we need to sequentially design some new points to query. This example can be viewed as a prediction-oriented active learning (AL) task (Smith et al., 2023).

The design space $\Xi=\mathcal{X}$ is the domain of x, and y is the corresponding noisy observations of the function. Let $\mathcal{Q}(\mathcal{X})$ denote the set of combinations of distributions that can be output by TNDP, we can then define decision space to be $\mathcal{A}=\mathcal{Q}(\mathcal{X})$. The downstream decision is to output a predictive distribution for y^* given a test point x^* , and the utility function $u(y_\Xi,a)=\log q(y^*|x^*,h_{1:t})$ is the log probability of y under the predicted distribution, given the queried historical data h_t .

During training, we sample functions from Gaussian Processes (GPs) (Rasmussen and Williams, 2006) with squared exponential kernels of varying output variances and lengthscales and randomly sample a point as the test point x^* . We set the global contextual information λ as the test point x^* . For illustration purposes, we consider only the case where T=1. Additional details for the data generation can be found in Appendix E.

Results. From Fig. 3(a), we can observe that the values of π concentrate near x^* , meaning our query head f_q tends to query points close to x^* to maximize the DUG. This is an intuitive example demonstrating that our TNDP can adjust its design strategy based on the downstream task.

6.2 Decision-aware active learning

We now show another application of our method in a case of decision-aware AL studied by Filstroff et al. (2024). In this experiment, the model will be used for downstream decision-making after performing AL, i.e., we will use the learned information to take an action towards a specific target. A practical application of this problem is personalized medical diagnosis introduced in Section 1, where a doctor needs to query historical patient data to decide on a treatment for a new patient.

We use the same problem setup as in Filstroff et al. (2024). The decision space consists of N_d available decisions, $a \in A \in \{1, ..., N_d\}$. The design space $\Xi = \mathcal{X} \times A$ is composed of the patient's

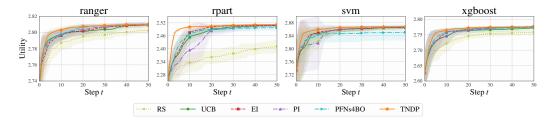


Figure 4: **Results on Top-***k* **HPO task.** For each meta-dataset, we calculated the average utility across all available test sets. The error bars represent the standard deviation over five runs. TNDP consistently achieved the best performance in terms of utility.

covariates x and the decisions a they receive. The outcome y is the treatment effect after the patient receives the decision, which is influenced by real-world unknown parameters θ such as medical conditions. For historical data, each patient is associated with only one decision. The utility function $u(y_{\Xi},a)=\mathbb{I}(\hat{a}^*,a^*)$ is a binary accuracy score that measures whether we can make the correct decision for a new patient x^* based on the queried history, where \hat{a}^* is the predicted Bayesian optimal decision and a^* the true optimal decision. Here, $u(y_{\Xi},a)=1$ if and only if $\hat{a}^*=a^*$.

In our experiment, we use the synthetic dataset from Filstroff et al. (2024), the details of the data generating process can be found in Appendix F. We set $N_d=4$ and use independent GPs to generate different outcomes. Each data point is randomly assigned a decision, and the outcome is the corresponding y value from the associated GP. We randomly select a test point x^* and determine the optimal decision a^* based on the GP that provides the maximum y value at x^* . We set global contextual information λ as the covariates of the test point x^* .

We compare TNDP with other non-amortized AL methods: random sampling (GP-RS), uncertainty sampling (GP-US), decision uncertainty sampling (GP-DUS), targeted information (GP-TEIG) introduced by Sundin et al. (2018), and decision EIG (GP-DEIG) proposed by Filstroff et al. (2024). A detailed description of each method can be found in Appendix F. Each method is tested on 100 different x^* points, and the average utility, i.e., the proportion of correct decisions, is calculated.

Results. The results are shown in Fig. 3(b), where we can see that TNDP achieves significantly better average accuracy than other methods. Additionally, we conduct an ablation study of TNDP in Appendix F.3 to verify the effectiveness of f_q . We further analyze the deployment running time to show the advantage of amortization, see Appendix D.1.

6.3 Top-k hyperparameter optimization

In traditional optimization tasks, we typically only aim to find a single point that maximizes the underlying function f. However, instead of identifying a single optimal point, there are scenarios where we wish to estimate a set of top-k distinct optima. For example, this is the case in robust optimization, where selecting multiple points can safeguard against variations in data or model performance.

In this experiment, we choose hyperparameter optimization (HPO) as our task and conduct experiments on the HPO-B datasets (Arango et al., 2021). The design space $\Xi \subseteq \mathcal{X}$ is a finite set defined over the hyperparameter space and the outcome y is the accuracy of a given configuration on a specific dataset. Our decision is to find k hyperparameter sets, denoted as $a=(a_1,...,a_k)\in A\subseteq \mathcal{X}^k$, with $a_i\neq a_j$. The utility function is then defined as $u(y_\Xi,a)=\sum_{i=1}^k y_{a_i}$, where y_{a_i} is the accuracy corresponding to the hyperparameter configuration a_i . In this experiment, the global contextual information $\lambda=\emptyset$.

We compare our methods with five different BO methods: random sampling (RS), Upper Confidence Bound (UCB), Expected Improvement (EI), Probability of Improvement (PI), and an amortized method PFNs4BO (Müller et al., 2023), which is a transformer-based model designed for hyperparameter optimization. We set k=3 and T=50, starting with an initial dataset of 5 points. Our experiments are conducted on four search spaces selected from the HPO-B benchmark. All results

are evaluated on a predefined test set, ensuring that TNDP does not encounter these test sets during training. For more details, see Appendix G.

Results. From the experimental results (Fig. 4), our method demonstrates superior performance across all four meta-datasets, particularly during the first 10 queries, achieving clearly better utility gains. This indicates that our TNDP can effectively identify high-performing hyperparameter configurations early in the optimization process.

Finally, we included a real-world experiment on retrosynthesis planning. Specifically, our task is to assist chemists in identifying the top-k synthetic routes for a novel molecule, as selecting the most practical routes from many random routes generated by the retrosynthesis software can be troublesome. The detailed description of the task and the results are shown in Appendix G.3.

7 Discussion

7.1 Limitations & future work

We recognize that the training of the query head inherently poses a reinforcement learning (RL) (Li, 2017) problem. Currently, we employ a basic REINFORCE algorithm, which can result in unstable training, particularly for tasks with sparse reward signals. For more complex problems in the future, we could deploy advanced RL methods, such as Proximal Policy Optimization (PPO) (Schulman et al., 2017); the trade-offs include the introduction of additional hyperparameters and increased computational cost. Like all amortized approaches, our method requires a large amount of data and upfront training time to develop a reliable model. Besides, our architecture is based on the Transformer, which suffers from quadratic complexity with respect to the input sequence length. This can become a bottleneck when the query set is very large. Future work could focus on designing more sample-efficient methods to reduce the data and training requirements. Our TNDP follows the common practice in the neural processes literature (Garnelo et al., 2018) of using independent Gaussian likelihoods. If modeling correlations between points is crucial for the downstream task, we can replace the output with a joint multivariate normal distribution (Markou et al., 2022) or predict the output autoregressively (Bruinsma et al., 2023). Following most BED approaches, our work assumes that the model is well-specified. However, model misspecification or shifts in the utility function during deployment could impact the performance of the amortized model (Rainforth et al., 2024). Future work could address the challenge of robust experimental design under model misspecification (Huang et al., 2023a). Another limitation is that our system is currently constrained to accepting designs of the same dimensionality. Future work could focus on developing dimension-agnostic methods to expand the scope of amortization. Lastly, our model is trained on a fixed-step length, assuming a finite horizon for the experimental design process. Future research could explore the design of systems that can handle infinite horizon cases, potentially improving the applicability of TNDP to a broader range of real-world problems.

7.2 Conclusions

In this paper, we proposed an amortized framework for decision-aware Bayesian experimental design (BED). We introduced the concept of *Decision Utility Gain* (DUG) to guide experimental design more effectively toward optimizing decision outcomes. Towards amortization, we developed a novel *Transformer Neural Decision Process* (TNDP) architecture with dual output heads: one for proposing new experimental designs and another for approximating the predictive distribution to facilitate optimal decision-making. Our experimental results demonstrated that TNDP significantly outperforms traditional BED methods across a variety of tasks. By integrating decision-making considerations directly into the experimental design process, TNDP not only accelerates the design of experiments but also improves the quality of the decisions derived from these experiments.

Acknowledgements

DH, LA and SK were supported by the Research Council of Finland (Flagship programme: Finnish Center for Artificial Intelligence FCAI). YG was supported by Academy of Finland grant 345604. LA was also supported by Research Council of Finland grants 358980 and 356498. SK was also supported by the UKRI Turing AI World-Leading Researcher Fellowship, [EP/W002973/1]. The authors wish to thank Aalto Science-IT project, and CSC–IT Center for Science, Finland, for the computational and data storage resources provided.

References

- Abbasnejad, E., Domke, J., and Sanner, S. (2015). Loss-calibrated monte carlo action selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.
- Arango, S. P., Jomaa, H. S., Wistuba, M., and Grabocka, J. (2021). Hpo-b: A large-scale reproducible benchmark for black-box hpo based on openml. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. (2020). Botorch: A framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33.
- Berger, J. O. (2013). Statistical decision theory and Bayesian analysis. Springer Science & Business Media.
- Bica, I., Alaa, A. M., Lambert, C., and Van Der Schaar, M. (2021). From real-world patient data to individualized treatment effects using machine learning: current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics*, 109(1):87–100.
- Bille, P. (2005). A survey on tree edit distance and related problems. *Theoretical computer science*, 337(1-3):217–239.
- Blacker, A. J., Williams, M. T., and Williams, M. T. (2011). *Pharmaceutical process development:* current chemical and engineering challenges, volume 9. Royal Society of Chemistry.
- Blau, T., Bonilla, E., Chades, I., and Dezfouli, A. (2023). Cross-entropy estimators for sequential experiment design with reinforcement learning. *arXiv* preprint arXiv:2305.18435.
- Blau, T., Bonilla, E. V., Chades, I., and Dezfouli, A. (2022). Optimizing sequential experimental design with deep reinforcement learning. In *International conference on machine learning*, pages 2107–2128, PMLR.
- Bruinsma, W., Markou, S., Requeima, J., Foong, A. Y., Andersson, T., Vaughan, A., Buonomo, A., Hosking, S., and Turner, R. E. (2023). Autoregressive conditional neural processes. In *The Eleventh International Conference on Learning Representations*.
- Burger, M., Hauptmann, A., Helin, T., Hyvönen, N., and Puska, J.-P. (2021). Sequentially optimized projections in x-ray imaging. *Inverse Problems*, 37(7):075006.
- Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical science*, pages 273–304.
- Chang, P. E., Loka, N., Huang, D., Remes, U., Kaski, S., and Acerbi, L. (2024). Amortized probabilistic conditioning for optimization, simulation and inference. *arXiv preprint arXiv:2410.15320*.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., and Mordatch, I. (2021). Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34.
- Cheng, Y. and Shen, Y. (2005). Bayesian adaptive designs for clinical trials. *Biometrika*, 92(3):633–646.
- Cobb, A. D., Roberts, S. J., and Gal, Y. (2018). Loss-calibrated approximate inference in bayesian neural networks. *arXiv preprint arXiv:1805.03901*.
- Filstroff, L., Sundin, I., Mikkola, P., Tiulpin, A., Kylmäoja, J., and Kaski, S. (2024). Targeted active learning for bayesian decision-making. *Transactions on Machine Learning Research*.
- Foster, A., Ivanova, D. R., Malik, I., and Rainforth, T. (2021). Deep adaptive design: Amortizing sequential bayesian experimental design. In *International Conference on Machine Learning*, pages 3384–3395. PMLR.
- Foster, A., Jankowiak, M., Bingham, E., Horsfall, P., Teh, Y. W., Rainforth, T., and Goodman, N. (2019). Variational bayesian optimal experimental design. *Advances in Neural Information Processing Systems*, 32.

- Foster, A., Jankowiak, M., O'Meara, M., Teh, Y. W., and Rainforth, T. (2020). A unified stochastic gradient approach to designing bayesian-optimal experiments. In *International Conference on Artificial Intelligence and Statistics*, pages 2959–2969. PMLR.
- Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D., and Eslami, S. A. (2018). Conditional neural processes. In *International conference on machine learning*, pages 1704–1713. PMLR.
- Garnett, R. (2023). Bayesian optimization. Cambridge University Press.
- Huang, D., Bharti, A., Souza, A., Acerbi, L., and Kaski, S. (2023a). Learning robust statistics for simulation-based inference under model misspecification. Advances in Neural Information Processing Systems, 36.
- Huang, D., Haussmann, M., Remes, U., John, S., Clarté, G., Luck, K., Kaski, S., and Acerbi, L. (2023b). Practical equivariances via relational conditional neural processes. Advances in Neural Information Processing Systems, 36.
- Ivanova, D. R., Foster, A., Kleinegesse, S., Gutmann, M. U., and Rainforth, T. (2021). Implicit deep adaptive design: Policy-based experimental design without likelihoods. *Advances in Neural Information Processing Systems*, 34.
- Ivanova, D. R., Hedman, M., Guan, C., and Rainforth, T. (2024). Step-DAD: Semi-Amortized Policy-Based Bayesian Experimental Design. ICLR 2024 Workshop on Data-centric Machine Learning Research (DMLR).
- Ivanova, D. R., Jennings, J., Rainforth, T., Zhang, C., and Foster, A. (2023). Co-bed: information-theoretic contextual optimization via bayesian experimental design. In *International Conference on Machine Learning*. PMLR.
- Kleinegesse, S. and Gutmann, M. U. (2020). Bayesian experimental design for implicit models by mutual information neural estimation. In *International conference on machine learning*, pages 5316–5326. PMLR.
- Kuśmierczyk, T., Sakaya, J., and Klami, A. (2019). Variational bayesian decision-making for continuous utilities. *Advances in Neural Information Processing Systems*, 32.
- Lacoste-Julien, S., Huszár, F., and Ghahramani, Z. (2011). Approximate inference for the loss-calibrated bayesian. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 416–424. JMLR Workshop and Conference Proceedings.
- Li, Y. (2017). Deep reinforcement learning: An overview. arXiv preprint arXiv:1701.07274.
- Lim, V., Novoseller, E., Ichnowski, J., Huang, H., and Goldberg, K. (2022). Policy-based bayesian experimental design for non-differentiable implicit models. *arXiv preprint arXiv:2203.04272*.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005.
- Lindley, D. V. (1972). Bayesian statistics: A review. SIAM.
- Maraval, A., Zimmer, M., Grosnit, A., and Bou Ammar, H. (2024). End-to-end meta-bayesian optimisation with transformer neural processes. *Advances in Neural Information Processing Systems*, 36.
- Markou, S., Requeima, J., Bruinsma, W., Vaughan, A., and Turner, R. E. (2022). Practical conditional neural process via tractable dependent predictions. In *International Conference on Learning Representations*.
- Mo, Y., Guan, Y., Verma, P., Guo, J., Fortunato, M. E., Lu, Z., Coley, C. W., and Jensen, K. F. (2021). Evaluating and clustering retrosynthesis pathways with learned strategy. *Chemical science*, 12(4):1469–1478.
- Morais, M. J. and Pillow, J. W. (2022). Loss-calibrated expectation propagation for approximate bayesian decision-making. *arXiv preprint arXiv:2201.03128*.

- Müller, S., Feurer, M., Hollmann, N., and Hutter, F. (2023). Pfns4bo: In-context learning for bayesian optimization. In *International Conference on Machine Learning*, pages 25444–25470. PMLR.
- Müller, S., Hollmann, N., Arango, S. P., Grabocka, J., and Hutter, F. (2021). Transformers can do bayesian inference. In *International Conference on Learning Representations*.
- Neiswanger, W., Wang, K. A., and Ermon, S. (2021). Bayesian algorithm execution: Estimating computable properties of black-box functions using mutual information. In *International Conference on Machine Learning*, pages 8005–8015. PMLR.
- Neiswanger, W., Yu, L., Zhao, S., Meng, C., and Ermon, S. (2022). Generalizing bayesian optimization with decision-theoretic entropies. Advances in Neural Information Processing Systems, 35.
- Nguyen, T. and Grover, A. (2022). Transformer neural processes: Uncertainty-aware meta learning via sequence modeling. In *International Conference on Machine Learning*, pages 16569–16594. PMLR.
- Parzen, E. (1999). Stochastic processes. SIAM.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Rainforth, T., Cornish, R., Yang, H., Warrington, A., and Wood, F. (2018). On nesting monte carlo estimators. In *International Conference on Machine Learning*, pages 4267–4276. PMLR.
- Rainforth, T., Foster, A., Ivanova, D. R., and Bickford Smith, F. (2024). Modern bayesian experimental design. Statistical Science, 39(1):100–114.
- Rasmussen, C. E. and Williams, C. K. (2006). Gaussian Processes for Machine Learning. MIT Press.
- Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. (2016). A review of modern computational algorithms for bayesian optimal design. *International Statistical Review*, 84(1):128– 154.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Smith, F. B., Kirsch, A., Farquhar, S., Gal, Y., Foster, A., and Rainforth, T. (2023). Prediction-oriented bayesian active learning. In *International Conference on Artificial Intelligence and Statistics*, pages 7331–7348. PMLR.
- Stevens, S. J. (2011). *Progress toward the synthesis of providencin*. PhD thesis, Colorado State University.
- Sundin, I., Peltola, T., Micallef, L., Afrabandpey, H., Soare, M., Mamun Majumder, M., Daee, P., He, C., Serim, B., Havulinna, A., et al. (2018). Improving genomics-based predictions for precision medicine through active elicitation of expert knowledge. *Bioinformatics*, 34(13):i395–i403.
- Szymkuć, S., Gajewska, E. P., Klucznik, T., Molga, K., Dittwald, P., Startek, M., Bajczyk, M., and Grzybowski, B. A. (2016). Computer-assisted synthetic planning: the end of the beginning. *Angewandte Chemie International Edition*, 55(20):5904–5937.
- Vadera, M. P., Ghosh, S., Ng, K., and Marlin, B. M. (2021). Post-hoc loss-calibration for bayesian neural networks. In *Uncertainty in Artificial Intelligence*, pages 1403–1412. PMLR.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, 30.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256.

- Zheng, Q., Zhang, A., and Grover, A. (2022). Online decision transformer. In *international conference* on machine learning. PMLR.
- Zhong, S., Shen, W., Catanach, T., and Huan, X. (2024). Goal-oriented bayesian optimal experimental design for nonlinear models using markov chain monte carlo. *arXiv preprint arXiv:2403.18072*.

Appendix

The appendix is organized as follows:

- In Appendix A, we provide a brief introduction to conditional neural processes (CNPs) and Transformer neural processes (TNPs).
- In Appendix B, we describe the details of our model architecture and the training setups.
- In Appendix C, we present the full algorithm for training our TNDP architecture.
- In Appendix D, we compare the inference time with other methods and show the overall training time of TNDP.
- In Appendix E, we describe the details of our toy example.
- In Appendix F, we describe the details of the decision-aware active learning example.
- In Appendix G, we describe the details of the top-k hyperparameter optimization task, along with additional results on the retrosynthesis planning task.

A Conditional neural processes

CNPs (Garnelo et al., 2018) are designed to model complex stochastic processes through a flexible architecture that utilizes a *context set* and a *target set*. The context set consists of observed data points that the model uses to form its understanding, while the target set includes the points to be predicted. The traditional CNP architecture includes an encoder and a decoder. The encoder is a DeepSet architecture to ensure permutation invariance, it transforms each context point individually and then aggregates these transformations into a single representation that captures the overall context. The decoder then uses this representation to generate predictions for the target set, typically employing a Gaussian likelihood for approximation of the true predictive distributions. Due to the analytically tractable likelihood, CNPs can be efficiently trained through maximum likelihood estimation.

A.1 Transformer neural processes

Transformer Neural Processes (TNPs), introduced by Nguyen and Grover (2022), improve the flexibility and expressiveness of CNPs by incorporating the Transformer's attention mechanism (Vaswani et al., 2017). In TNPs, the transformer architecture uses self-attention to process the context set, dynamically weighting the importance of each point. This allows the model to create a rich representation of the context, which is then used by the decoder to generate predictions for the target set. The attention mechanism in TNPs facilitates the handling of large and variable-sized context sets, improving the model's performance on tasks with complex input-output relationships. The Transformer architecture is also useful in our setups where certain designs may have a more significant impact on the decision-making process than others. For more details about TNPs, please refer to Nguyen and Grover (2022).

B Implementation details

B.1 Embedders

The embedder $f_{\rm emb}$ is responsible for mapping the raw data to a space of the same dimension. For the toy example and the top-k hyperparameter task, we use three embedders: a design embedder $f_{\rm emb}^{(\xi)}$, an outcome embedder $f_{\rm emb}^{(y)}$, and a time step embedder $f_{\rm emb}^{(t)}$. Both $f_{\rm emb}^{(\xi)}$ and $f_{\rm emb}^{(y)}$ are multi-layer perceptions (MLPs) with the following architecture:

- **Hidden dimension**: the dimension of the hidden layers, set to 32.
- Output dimension: the dimension of the output space, set to 32.
- **Depth**: the number of layers in the neural network, set to 4.
- Activation function: ReLU is used as the activation function for the hidden layers.

The time step embedder $f_{\text{emb}}^{(t)}$ is a discrete embedding layer that maps time steps to a continuous embedding space of dimension 32.

For the decision-aware active learning task, since the design space contains both the covariates and the decision, we use four embedders: a covariate embedder $f_{\rm emb}^{(x)}$, a decision embedder $f_{\rm emb}^{(d)}$, an outcome embedder $f_{\rm emb}^{(y)}$, and a time step embedder $f_{\rm emb}^{(t)}$. $f_{\rm emb}^{(x)}$, $f_{\rm emb}^{(y)}$ and $f_{\rm emb}^{(t)}$ are MLPs which use the same settings as described above. The decision embedder $f_{\rm emb}^{(d)}$ is another discrete embedding layer.

For context embedding $E^{(c)}$, we first map each $\xi_i^{(c)}$ and $y_i^{(c)}$ to the same dimension using their respective embedders, and then sum them to obtain the final embedding. For prediction embedding $E^{(p)}$ and query embedding $E^{(q)}$, we only encode the designs. For E^{GI} , except the embeddings of the time step, we also encode the global contextual information λ using $f_{\rm emb}^{(x)}$ in the toy example and the decision-aware active learning task. All the embeddings are then concatenated together to form our final embedding E.

B.2 Transformer blocks

We utilize the official TransformerEncoder layer of PyTorch (Paszke et al., 2019) (https://pytorch.org) for our transformer architecture. For all experiments, we use the same configuration: the model has 6 Transformer layers, with 8 heads per layer, the MLP block has a hidden dimension of 128, and the embedding dimension size is set to 32.

B.3 Output heads

The prediction head, f_p is an MLP that maps the Transformer's output embeddings of the query set to the predicted outcomes. It consists of an input layer with 32 hidden units, a ReLU activation function, and an output layer. The output layer predicts the mean and variance of a Gaussian likelihood, similar to CNPs.

For the query head f_q , all candidate experimental designs are first mapped to embeddings $\lambda^{(q)}$ by the Transformer, and these embeddings are then passed through f_q to obtain individual outputs. We then apply a Softmax function to these outputs to ensure a proper probability distribution. f_q is an MLP consisting of an input layer with 32 hidden units, a ReLU activation function, and an output layer.

B.4 Training details

For all experiments, we use the same configuration to train our model. We set the initial learning rate to 5e-4 and employ the cosine annealing learning rate scheduler. The number of training epochs is set to 50,000 for top-k tasks and 100,000 for other tasks, and the batch size is 16. For the REINFORCE, we use a discount factor of $\alpha=0.99$.

C Full algorithm for training TNDP

```
Algorithm 1 Transformer Neural Decision Processes (TNDP)
```

```
1: Input: Utility function u(y_{\Xi}, a), prior p(\theta), likelihood p(y|\theta, \xi), query horizon T
 2: Output: Trained TNDP
     while within the training budget do
          Sample \theta \sim p(\theta) and initialize D
 4:
          for t=1 to T do \xi_t^{(q)} \sim \pi_t(\cdot|h_{1:t-1}) Sample y_t \sim p(y|\theta,\xi)
 5:
                                                                                                      ⊳ sample next design from policy
 7:
                                                                                                                             Set h_{1:t} = h_{1:t-1} \cup \{(\xi_t^{(q)}, y_t)\}

Set D^{(c)} = h_{1:t}, D^{(q)} = D^{(q)} \setminus \{\xi_t^{(q)}\}

Calculate r_t(\xi_t^{(q)}) with u(y_\Xi, a) using Eq. (9)
 8:

    □ update history

 9:
                                                                                                                                         \triangleright update D
10:

    ▷ calculate reward

         end for R_t = \sum_{k=t}^{T} \alpha^{k-t} r_k(\xi_k^{(q)}) Update TNDP using \mathcal{L}^{(p)} (Eq. (7)) and \mathcal{L}^{(q)} (Eq. (10))
11:
12:
                                                                                                           ⊳ calculate cumulative reward
13:
15: At deployment, we can use f^{(q)} to sequentially query T designs. Afterward, based on the queried
```

experiments, we perform one-step final decision-making using the prediction from $f^{(p)}$.

D Computational cost analysis

D.1 Inference time analysis

We evaluate the inference time of our algorithm during the deployment stage. We select decision-aware active learning as the experiment for our time comparison. All experiments are evaluated on an Intel Core i7-12700K CPU. We measure both the acquisition time and the total time. The acquisition time refers to the time required to compute one next design, while the total time refers to the time required to complete 10 rounds of design. The final results are presented in Table A1, with the mean and standard deviation calculated over 10 runs.

Traditional methods rely on updating the GP and optimizing the acquisition function, which is computationally expensive. D-EIG and T-EIG require many model retraining steps to get the next design, which is not tolerable in applications requiring fast decision-making. However, since our model is fully amortized, once it is trained, it only requires a single forward pass to design the experiments, resulting in significantly faster inference times.

Method	Acquisition time (s)	Total time (s)
GP-RS	0.00002(0.00001)	28(7)
GP-US	0.07(0.01)	29(7)
GP-DUS	0.38(0.02)	44(5)
T-EIG (Sundin et al., 2018)	1558(376)	15613(3627)
D-EIG (Filstroff et al., 2024)	572(105)	5746 (1002)
TDNP (ours)	0.015(0.004)	0.31(0.06)

Table A1: Comparison of computational costs across different methods. We report the mean value and (standard deviation) derived from 10 runs with different seeds.

D.2 Overall training time

Throughout this paper, we carried out all experiments, including baseline model computations and preliminary experiments not included in the final paper, on a GPU cluster featuring a combination of Tesla P100, Tesla V100, and Tesla A100 GPUs. We estimate the total computational usage to be roughly 5000 GPU hours. For each experiment, it takes around 10 GPU hours on a Tesla V100 GPU with 32GB memory to reproduce the result, with an average memory consumption of 8 GB.

E Details of toy example

E.1 Data generation

In our toy example, we generate data using a GP with the Squared Exponential (SE) kernel, which is defined as:

$$k(x, x') = v \exp\left(-\frac{(x - x')^2}{2\ell^2}\right),\tag{A1}$$

where v is the variance, and ℓ is the lengthscale. Specifically, in each training iteration, we draw a random lengthscale $\ell \sim 0.25 + 0.75 \times U(0,1)$ and the variance $v \sim 0.1 + U(0,1)$, where U(0,1) denotes a uniform random variable between 0 and 1.

F Details of decision-aware active learning experiments

F.1 Data generation

For this experiment, we use a GP with a Squared Exponential (SE) kernel to generate our data. The covariates x are drawn from a standard normal distribution. For each decision, we use an independent GP to simulate different outcomes. In each training iteration, the lengthscale for each GP is randomly sampled as $\ell \sim 0.25 + 0.75 \times U(0,1)$ and the variance as $v \sim 0.1 + U(0,1)$, where U(0,1) denotes a uniform random variable between 0 and 1.

F.2 Other methods description

We compare our method with other non-amortized approaches, all of which use GPs as the functional prior. Each model is equipped with an SE kernel with automatic relevance determination. GP hyperparameters are estimated with maximum marginal likelihood.

Our method is compared with the following methods:

- Random sampling (GP-RS): randomly choose the next design ξ_t from the query set.
- Uncertainty sampling (GP-US): choose the next design ξ_t for which the predictive distribution $p(y_t|\xi_t,h_{t-1})$ has the largest variance.
- Decision uncertainty sampling (GP-DUS): choose the next design ξ_t such that the predictive distribution of the optimal decision corresponding to this design is the most uncertain.
- Targeted information (GP-TEIG) (Sundin et al., 2018): a targeted active learning criterion, introduced by (Sundin et al., 2018), selects the next design ξ_t that provides the highest EIG on $p(y^*|x^*, h_{t-1} \cup \{(\xi_t, y_t)\})$.
- Decision EIG (GP-DEIG) (Filstroff et al., 2024): choose the next design ξ_t which directly aims at reducing the uncertainty on the posterior distribution of the optimal decision. See Filstroff et al. (2024) for a detailed explanation.

F.3 Ablation study

We also carry out an ablation study to verify the effectiveness of our query head and the non-myopic objective function. We first compare TNDP with TNDP using random sampling (TNDP-RS), and the results are shown in Fig. A1(a). We observe that the designs proposed by the query head significantly improve accuracy, demonstrating that the query head can propose informative designs based on downstream decisions.

We also evaluate the impact of the non-myopic objective by comparing TNDP with a myopic version that only optimizes for immediate utility rather than long-term gains ($\alpha=0$). The results, presented in Fig. A1(b), show that TNDP with the non-myopic objective function achieves higher accuracy across iterations compared to using the myopic objective. This indicates that our non-myopic objective effectively captures the long-term benefits of each design choice, leading to improved overall performance.

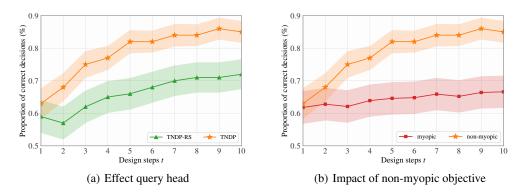


Figure A1: **Comparison of TNDP variants on the decision-aware active learning task.** (a) Shows the effect of the query head, where TNDP outperforms TNDP-RS, demonstrating its ability to generate informative designs. (b) Illustrates the impact of the non-myopic objective, with TNDP achieving higher accuracy than the myopic version.

G Details of top-k hyperparameter optimization experiments

G.1 Data

In this task, we use HPO-B benchmark datasets (Arango et al., 2021). The HPO-B dataset is a large-scale benchmark for HPO tasks, derived from the OpenML repository. It consists of 176 search spaces (algorithms) evaluated on 196 datasets, with a total of 6.4 million hyperparameter evaluations. This dataset is designed to facilitate reproducible and fair comparisons of HPO methods by providing explicit experimental protocols, splits, and evaluation measures.

We extract four meta-datasets from the HPOB dataset: ranger (id=7609, d_x =9), svm (id=5891, d_x =8), rpart (id=5859, d_x =6), and xgboost (id=5971, d_x =16). In the test stage, the initial context set is chosen based on their pre-defined indices. For detailed information on the datasets, please refer to https://github.com/releaunifreiburg/HPO-B.

G.2 Other methods description

In our experiments, we compare our method with several common acquisition functions used in HPO. We use GPs as surrogate models for these acquisition functions. All the implementations are based on BoTorch (Balandat et al., 2020) (https://botorch.org/). The acquisition functions compared are as follows:

- Random Sampling (RS): This method selects hyperparameters randomly from the search space, without using any surrogate model or acquisition function.
- Upper Confidence Bound (UCB): This acquisition function balances exploration and exploitation by selecting points that maximize the upper confidence bound. The UCB is defined as:

$$\alpha_{\text{UCB}}(\mathbf{x}) = \mu(\mathbf{x}) + \kappa \sigma(\mathbf{x}),$$
 (A2)

where $\mu(\mathbf{x})$ is the predicted mean, $\sigma(\mathbf{x})$ is the predicted standard deviation, and κ is a parameter that controls the trade-off between exploration and exploitation.

• Expected Improvement (EI): This acquisition function selects points that are expected to yield the greatest improvement over the current best observation. The EI is defined as:

$$\alpha_{\text{EI}}(\mathbf{x}) = \mathbb{E}[\max(0, f(\mathbf{x}) - f(\mathbf{x}^+))],\tag{A3}$$

where $f(\mathbf{x}^+)$ is the current best value observed, and the expectation is taken over the predictive distribution of $f(\mathbf{x})$.

• **Probability of Improvement (PI)**: This acquisition function selects points that have the highest probability of improving over the current best observation. The PI is defined as:

$$\alpha_{\text{PI}}(\mathbf{x}) = \Phi\left(\frac{\mu(\mathbf{x}) - f(\mathbf{x}^+) - \omega}{\sigma(\mathbf{x})}\right),$$
(A4)

where Φ is the cumulative distribution function of the standard normal distribution, $f(\mathbf{x}^+)$ is the current best value observed, and ω is a parameter that encourages exploration.

In addition to those non-amortized GP-based methods, we also compare our method with an amortized surrogate model **PFNs4BO** (Müller et al., 2023). It is a Transformer-based model designed for hyperparameter optimization which does not consider the downstream task. We use the pre-trained PFNs4BO-BNN model which is trained on HPO-B datasets and choose PI as the acquisition function, the model and the training details can be found in their official implementation (https://github.com/automl/PFNs4BO).

G.3 Additional experiment on retrosynthesis planning

We now show a real-world experiment on retrosynthesis planning (Blacker et al., 2011). Specifically, our task is to help chemists identify the top-k synthetic routes for a novel molecule (Mo et al., 2021), as it can be challenging to select the most practical routes from many random options generated by the retrosynthesis software (Stevens, 2011; Szymkuć et al., 2016). In this task, the design space for each molecule m is a finite set of routes that can synthesize the molecule. The sequential experimental

design is to select a route for a specific molecule to query its score y, which is calculated based on the tree edit distance (Bille, 2005) from the best route. The downstream task is to recommend the top-k routes with the highest target-specific scores based on the collected information.

In this experiment, we choose k = 3 and T=10, and set the global information $\gamma=m$. We train our TNDP on a novel non-public metadataset, including 1500 molecules with 70 synthetic routes for each molecule. The representation dimension of the molecule is 64 and that of the route is 264, both of which are learned through a neural network. Given the highdimensional nature of the data representation, it is challenging to directly compare TNDP with other GP-based methods, as GPs typically struggle with scalability and performance in such high-dimensional settings. Therefore, we only compare TNDP with TNDP using random sampling. The final results are evaluated on 50 test molecules that are not seen during training, as shown in Fig. A2.

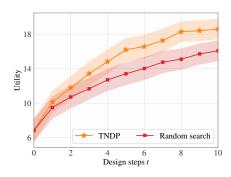


Figure A2: **Results of retrosynthesis planning experiment.** The utility is the sum of the quality scores of top-k routes and is calculated with 50 molecules. Our TNDP outperforms the random search baseline.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: In this paper, we propose an amortized decision-aware BED framework (Section 3 and Section 4), addressing the desiderate delineated in the abstract and in the introduction. We explicitly state the contributions at the end of the introduction. The experimental results presented in Section 6 match the claims made.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations of our work explicitly in Section 7.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We fully describe our model architecture in Section 4, and we also provide the implementation details in Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We have included the code repository link in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We describe the experimental settings for all experiments in Appendix E, Appendix F, Appendix G.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All the experiment results are accompanied by error bars. We also report the method for calculating the error bars in each experiment section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We report the compute resources in Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: Our work adheres to the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The work is foundational research and not tied to particular applications. We believe there is no negative societal impact on the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not pose such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original paper that produced the dataset, and also provide the URLs for the code packages we use.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.