
ESCADA: Efficient Safety and Context Aware Dose Allocation for Precision Medicine

Ilker Demirel*
Bilkent University
ilkerd@ee.bilkent.edu.tr

A. Alparslan Celik
Bilkent University
acelik@ee.bilkent.edu.tr

Cem Tekin
Bilkent University
cemtekin@ee.bilkent.edu.tr

Abstract

Finding an optimal individualized treatment regimen is considered one of the most challenging precision medicine problems. Various patient characteristics influence the response to the treatment, and hence, there is no one-size-fits-all regimen. Moreover, the administration of an unsafe dose during the treatment can have adverse effects on health. Therefore, a treatment model must ensure patient *safety* while *efficiently* optimizing the course of therapy. We study a prevalent medical problem where the treatment aims to keep a physiological variable in a safe range and preferably close to a target level, which we refer to as *leveling*. Such a task may be relevant in numerous other domains as well. We propose ESCADA, a novel and generic multi-armed bandit (MAB) algorithm tailored for the leveling task, to make safe, personalized, and context-aware dose recommendations. We derive high probability upper bounds on its cumulative regret and safety guarantees. Following ESCADA’s design, we also describe its Thompson sampling-based counterpart. We discuss why the straightforward adaptations of the classical MAB algorithms such as GP-UCB may not be a good fit for the leveling task. Finally, we make *in silico* experiments on the bolus-insulin dose allocation problem in type-1 diabetes mellitus disease and compare our algorithms against the famous GP-UCB algorithm, the rule-based dose calculators, and a clinician.

1 Introduction

Precision medicine aims to provide the best possible treatment on an individual level by considering patient characteristics’ variability [3, 30]. Many healthcare problems require keeping a physiological variable (e.g., blood glucose level) in a *safe* range and preferably close to a target level. One such example is electrolyte disorders, common among intensive care unit patients. When the blood sodium level falls below 135 milliequivalents per liter (mEq/L) or goes beyond 145 mEq/L, the patient experiences hypo-/hyper-natremia with adverse effects on health [24]. Therefore, correct dosing of electrolytes is crucial to ensure patient safety, and there is no consensus on how to assess the correct dosage for different patient characteristics. Another critical problem is blood pressure disorder. These are hypo-/hyper-tension events where the blood pressure deviates from its standard value and needs to be corrected. Patient characteristics play an essential role in determining the blood pressure response to the therapeutic agent, and they should be taken into account in the dosing process [33].

Related work and background A fair amount of research is dedicated to adaptive clinical trials which aim to identify a drug’s effectiveness within a group, including a tradeoff between efficacy and

*Now a graduate student at MIT CSAIL. Email: demirel@mit.edu

toxicity [26, 27, 40, 53]. The algorithms proposed in these works are not applicable to the problem structure considered here for two main reasons. First, the therapeutic agent is not necessarily *toxic*, and our aim is not to maximize the response to the agent but to keep it close to a target level. Therefore, classical upper confidence bound (UCB) based algorithms such as UCB1 [4] or GP-UCB [41] are not applicable for our objective. That is simply because the UCB-based algorithms leverage the *optimism in the face of uncertainty* (OFU) principle to form optimistic estimates of arm outcomes and pick the arm with the highest estimated outcome. However, in our case, *optimism* refers to the proximity of an arm’s outcome to the target. This fundamental difference in our task necessitates a novel acquisition strategy. One could simply form pseudo-rewards to maximize, such as $r(n) = -|o(n) - K|$, where $o(n)$ is the outcome at the end of round n and K is the target level. We particularly refrain from doing so as different reasonable choices for the pseudo-reward will lead the algorithm to operate differently in practice. Therefore, we keep the objective (i.e., minimize $|o(n) - K|$) in the most generic form and propose a suitable acquisition strategy instead. We provide more details on our objective and motivation behind designing a new acquisition strategy in §2. Secondly, our goal is to provide personalized recommendations rather than for a group of patients. We approach the safe dose allocation problem from a contextual multi-armed bandit (MAB) [29] perspective with additional safety constraints and propose a novel acquisition function tailored for this problem structure in §3.

To render our acquisition method safe, we propose a safe exploration strategy. There is a surge of interest in safe exploration for Bayesian optimization (BO), Markov decision processes, MABs, and reinforcement learning in general. [15, 17, 31, 54]. [1] propose the linear Thompson sampling (LTS) algorithm for the linear stochastic bandit (LSB) setting by adding a random perturbation to the regularized least-squares estimates of the parameters in a way that the OFU principle can be used. [32] modifies the LTS’ randomization procedure to continue leveraging the OFU principle in the face of additional safety constraints and matches LTS’ order of regret. [21] proposes a safe algorithm incurring a near-optimal expected regret for the LSB problem as well, which uses the arm outcomes’ lower confidence bounds to guarantee the safety of exploration and greedily exploit when it is safe.

There is a strand of literature on “risk-averse” MABs, where the learner is concerned not only with maximizing long-term earnings but also with reducing a certain measure of *risk* [9, 37]. [37, 49, 50] investigate the MAB problem using two risk measures, Mean-Variance and Value-at-Risk, which are widely adopted in financial portfolio management [42]. [8, 14, 20] study the Conditional-Value-at-Risk measure, which captures the tail-risk better compared to the Value-at-Risk measure. Another related area is the “conservative” bandits, where the learner’s cumulative reward must always exceed a predetermined fraction of a baseline’s [19, 56]. These works, however, do not address *stagewise* safety constraints on instantaneous arm outcomes, which must be explicitly satisfied at any given time.

We operate in a BO framework where we model the objective function as a sample from a Gaussian process (GP). [15, 17] consider BO with stagewise safety constraints. However, they aim to find optimal safe solutions and allow unsafe evaluations during exploration. [2] propose a safe variant of GP-UCB, which employs a pure exploration phase at the beginning and provides upper bounds on its cumulative regret. SafeOPT and StageOPT algorithms provide guarantees on the safety of the exploration process [44, 45]. However, they model the exploration of the safe set as a proxy objective which leads to unnecessary suboptimal evaluations at the boundaries of the safe set [48]. Moreover, they do not provide formal regret bounds. Goal-oriented Safe Expansion (GoOSE) algorithm works with any acquisition function as a *plug-in* safety mechanism and encourages the expansion of the safe set only when necessary [48]. When the query is not guaranteed to be safe, only then GoOSE expands the safe set by evaluating the function at safe points to learn more about the initial query’s safety. However, such re-evaluations are not possible within the framework of dynamic treatment regimes since this setup does not allow the administration of multiple doses. Moreover, all the works above consider a one-sided safety constraint ($f(x) \geq c$), whereas we consider a two-sided one as the aim is to keep $f(x)$ in a range ($c_1 \leq f(x) \leq c_2$). We provide a table comparing our work to some existing literature on safe exploration with GPs in Appendix A. Our key contributions are as follows.

We study an important and overlooked problem in medicine that which is relevant in other domains as well, such as demand-side management [7]. We formalize the problem from a contextual MAB perspective via a suitable definition of *regret* as the proxy performance metric in §2. Since our objective is to keep the outcomes close to a target rather than maximize them as in the usual MAB setting, we propose a novel acquisition method in §3. We design a safe exploration scheme for our acquisition function in §3 and derive high probability upper bounds on its regret with safety guarantees in §4. We make *in silico* experiments on type-1 diabetes mellitus (T1DM) disease in

§5. T1DM is characterized by insulin deficiency due to pancreatic β -cell loss, and it can have adverse effects which might result in hospitalization and death [6]. Therefore, T1DM patients must regulate their blood glucose by administering bolus insulin doses before meals. We optimize the dose recommendation process via *safely* and *efficiently* learning to recommend better doses.

2 Problem statement

We denote by $[N]$ the set $\{1, \dots, N\}$, $z \in \mathcal{Z}$ a context, and $d \in \mathcal{D}$ a dose, where both \mathcal{Z} and \mathcal{D} are compact and convex, and $\mathcal{D} = [0, \bar{D}]$. Let $f : \mathcal{Z} \times \mathcal{D} \rightarrow \Omega$ be the *unknown* function that maps (z, d) pairs to the physiological variable of interest, where $\Omega = [0, \bar{T}]$. At round $n \in [N]$, the learner observes a context, z_n , and recommends a dose, d_n , to obtain a noisy evaluation of f at (z_n, d_n) , given as $y_n = f(z_n, d_n) + \nu_n$, where ν_n are zero-mean i.i.d. Gaussian with known variance σ^2 . The learner's objective is to keep the physiological variable, $f(z_n, d_n)$, within a safe range and close to the target level. We formalize this objective as a contextual MAB problem with safety constraints as,

$$\text{minimize } R_N = \sum_{n=1}^N |f(z_n, d_n) - T| \quad (1)$$

$$\text{subject to } T_{\min} \leq f(z_n, d_n) \leq T_{\max}, \quad \forall n \in [N], \quad (2)$$

where T_{\min} and T_{\max} denote the lower and upper safety thresholds for f , respectively, and $T \in (T_{\min} + \alpha, T_{\max} - \alpha)$ is the target value, where $\alpha > 0$. We introduce the non-zero α term to ensure that the target level is not exactly equal to the safety thresholds, which is required later in the analyses. We assume $\forall z \in \mathcal{Z}$, there exists $d_z^* \in \mathcal{D}$ such that $f(z, d_z^*) = T$.




Figure 1: A hypothetical objective function $f(x)$, and three candidate reward functions.

Regularity assumptions Our safe exploration strategy relies on expanding around an initial safe set by exploiting the smoothness properties of the objective function $f(x)$. Without an initial safe set, and some regularity assumptions on $f(x)$, it is not possible to make inferences on the safety of the prospective recommendations [44]. Let $\mathcal{X} = \mathcal{Z} \times \mathcal{D}$ denote the space of all context-dose pairs. Let $k(\cdot, \cdot)$ be a positive definite kernel function on \mathcal{X} . We assume that $f(x)$ is a function from the *Reproducing Kernel Hilbert Space* (RKHS) corresponding to $k(\cdot, \cdot)$. In addition, we assume that $f(x)$ has bounded norm in this particular RKHS, i.e., $\|f\|_k < B_f$ [39]. This mild assumption makes $f(x)$ smooth enough to be efficiently learnable by a GP. More precisely, $f(x)$ is L -Lipschitz continuous w.r.t. kernel

metric $q(x, x') = \sqrt{k(x, x) - 2k(x, x') + k(x', x')}$, where $L = B_f$ [43]. Also, we denote by $q_z(d, d') := q((z, d), (z, d'))$. At this point, we define a discretization of \mathcal{D} for every $z \in \mathcal{Z}$ as,

$$\bar{\mathcal{D}}_z := \{d_i(z) \in \mathcal{D} \mid i \in \{1, \dots, k\}\},$$

where $d_1(z) = 0$, $d_i(z) > d_j(z)$ for $i > j$, $q_z(d_i, d_{i+1}) = \lambda/2L$, $q_z(d_k, \bar{D}) < \lambda/2L$, and $\lambda > 0$ is the discretization parameter. We assume that an initial safe set of discretized doses $S_0(z)$ is available for each $z \in \mathcal{Z}$. These assumptions allow us to use Gaussian processes (GP) to design our algorithm, and analyze its regret and safety guarantees [36]. A GP is a distribution over functions which is characterized by its mean, $\mu(\cdot)$, and covariance, $k(\cdot, \cdot)$, functions. Once we assume a GP prior over $f(x)$, after observing a set of noisy evaluations $\mathbf{y}_N = [y_1 \dots y_N]^T$ at $A_N = \{x_1, \dots, x_N\}$, the posterior over $f(x)$ is a GP again with the following mean and covariance functions,

$$\begin{aligned} k_N(x, x') &= k(x, x') - \mathbf{k}_N(x)^T (\mathbf{K}_N + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_N(x') \\ \sigma_N^2(x) &= k_N(x, x) \\ \mu_N(x) &= \mathbf{k}_N(x)^T (\mathbf{K}_N + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_N, \end{aligned}$$

where $\mathbf{k}_N(x) = [k(x_1, x), \dots, k(x_N, x)]^T$ and \mathbf{K}_N is the positive definite kernel matrix $[k(x, x')]_{x, x' \in A_N}$.

Comparison with GP-UCB Our objective is to keep $f(x)$ close to a target level T . As we discussed in §1, one could use the GP-UCB algorithm in [41] if the objective was to *maximize* $f(x)$. In our case, however, we have to define *pseudo-rewards* to maximize such as $-|f(z_n, d_n) - T|$ that are decreasing




Figure 2: *ESCADA Algorithm Description (left)*. Upon observing a context z_n in round n , TACO forms the set $C_n(z_n) \subseteq S_n(z_n)$ after eliminating the doses that are suboptimal with high probability (TACO uses $\mathcal{D}_n = S_n(z_n)$ in Algorithm 1 to ensure safety with high probability). If $C_n(z_n) \neq \emptyset$, it recommends the dose whose mean response is closest to the target T . If $C_n(z_n) = \emptyset$, it recommends the dose with the widest confidence interval in $S_n(z_n) \cap \bar{\mathcal{D}}_{z_n}$. *Flowchart (right)*. A simple interpretation of the dose allocation process intended for domain experts.

with $|f(z_n, d_n) - T|$ to capture the “leveling” task in (1). We present three such reward functions in Figure 1. However, $f(x)$ being smooth and efficiently learnable by a GP does not imply that a reward functional defined on $f(x)$ will be as well. Figure 1 shows that different reward functions can have significantly different landscapes. For instance, it is almost impossible to achieve our task efficiently by using the so-called “plausible” Reward 1. In §5, we compare our algorithms’ performances against the GP-UCB’s for three reward functions in Figure 1. Also, when $T \neq (T_{\min} + T_{\max})/2$ (which may well be the case, see §5), the reward-based GP-UCB method needs another GP to directly learn $f(x)$ to efficiently satisfy the safety requirements, doubling the computational complexity compared to our algorithms which use a single GP for everything. Finally, by learning $f(x)$ with a GP, we can provide interpretations for our model’s recommendations (see Figure 2).

3 ESCADA algorithm

We propose ESCADA: Efficient Safety and Context Aware Dose Allocation algorithm. It consists of two blocks: (i) an acquisition function, which we call TACO: Target-based Confident-acquisition, (ii) a safety mechanism to render TACO safe. Algorithm 1 and Figure 2 summarize ESCADA’s design. ESCADA’s recommendation procedure can be interpreted to domain experts via the flowchart in Figure 2 as opposed to black-box models [58].²

Acquisition strategy We propose TACO, a novel acquisition method specifically tailored for the “leveling” task described in §2. At each round n , TACO uses the confidence bounds of doses $d \in \mathcal{D}$ for z_n derived from the GP prior as $l_n(z_n, d) = \mu_{n-1}(z_n, d) - \beta_n^{1/2} \sigma_{n-1}(z_n, d)$, and $u_n(z_n, d) = \mu_{n-1}(z_n, d) + \beta_n^{1/2} \sigma_{n-1}(z_n, d)$. We define β_n later in a way that the confidence intervals contain the true value of f with high probability (see Lemma 1). Then, using Lipschitz continuity of f , we form the final lower and upper confidence bounds for every $d \in \mathcal{D}$ as,

$$\bar{l}_n(z_n, d) = \max\{l_n(z_n, d), l_n(z_n, d') - Lq_{z_n}(d, d')\}$$

$$\bar{u}_n(z_n, d) = \min\{u_n(z_n, d), u_n(z_n, d') + Lq_{z_n}(d, d')\},$$

where $d' = \arg\min_{\hat{d} \in \bar{\mathcal{D}}_{z_n}} q_{z_n}(d, \hat{d})$. We denote by $C_n(z_n, d) = [\bar{l}_n(z_n, d), \bar{u}_n(z_n, d)]$ the confidence interval of a dose $d \in \mathcal{D}$ in round n , and by $C_n(z_n) = \{C_n(z_n, d)\}_{d \in \mathcal{D}}$. Finally, we form the confidence widths for each dose $d \in \mathcal{D}$ as $w_n(z_n, d) = \bar{u}_n(z_n, d) - \bar{l}_n(z_n, d)$.

TACO queries a recommendation from a dose set \mathcal{D}_n at each round n upon observing the context z_n in three steps: (i): Identify the dose set $C_n \subseteq \mathcal{D}_n$ whose elements’ confidence intervals contain

Algorithm 1 ESCADA algorithm

for $n = 1, 2, \dots$ **do**
 Observe z_n and form $C_n(z_n)$
 Update $S_n(z_n)$ via (3)
 $d_n \leftarrow \text{TACO}(C_n(z_n), S_n(z_n))$
 Observe $y_n = f(z_n, d_n) + \nu_n$
 Update GP posterior

Subroutine: TACO

Inputs: $C_n(z_n); \mathcal{D}_n$
 $C_n = \{d \in \mathcal{D}_n \mid T \in C_n(z_n, d)\}$
if $C_n \neq \emptyset$ **then**
 $d \leftarrow \arg\min_{d' \in C_n} |\mu_{n-1}(z_n, d') - T|$
else
 $d \leftarrow \arg\max_{d' \in \mathcal{D}_n \cap \bar{\mathcal{D}}_{z_n}} w_n(z_n, d')$
return d

²Flowchart assumes that GP-induced confidence intervals are correct, i.e., the event \mathcal{E} in Lemma 1 holds.

the target value, T . (ii) If $\mathcal{C}_n \neq \emptyset$, recommend the dose in \mathcal{C}_n with the closest mean response to the target value T . (iii) If $\mathcal{C}_n = \emptyset$, recommend the dose in $\mathcal{D}_n \cap \overline{\mathcal{D}}_{z_n}$ with the widest confidence interval. In the first step, TACO eliminates the doses which are suboptimal with high probability. This step includes elements of both *exploration* and *exploitation*. A dose whose mean response is close to the target value can be selected (exploitation). On the other hand, if a dose is under-explored, it will have a wider confidence interval which may contain the target, and it stands a chance to be selected (exploration). In the third step, TACO focuses on pure exploration to identify the doses that may be optimal. TACO is *efficient* in the sense that it treats exploration as a proxy objective –in the third step– only when all the feasible doses (i.e., safe) are suboptimal with high probability.

Safety awareness We design a safe exploration scheme inspired from the previous works on safe GP optimization [44, 45]. We denote the safe set at round n for the context z_n by $S_n(z_n)$. Let us denote by $\hat{l}_n(z_n, d, d') := \bar{l}_n(z_n, d) - Lq_{z_n}(d, d')$, and $\hat{u}_n(z_n, d, d') := \bar{u}_n(z_n, d) + Lq_{z_n}(d, d')$. We implement the following expansion rule to derive $S_n(z_n)$ each round,

$$S_n(z_n) = S_{n-1}(z_n) \cup \left(\bigcup_{d \in S_{n-1}(z_n)} \{d' \in \mathcal{D} \mid \hat{l}_n(z_n, d, d') \geq T_{\min} \wedge \hat{u}_n(z_n, d, d') \leq T_{\max}\} \right), \quad (3)$$

To satisfy the safety requirements, TACO recommends a dose from $\mathcal{D}_n = S_n(z_n)$ at each round n . $S_n(z_n)$ only contains the doses for which f resides in the target interval almost certainly (see Theorem 1). We also define the ϵ -reachability operator \mathcal{R}_ϵ , where $\epsilon > 0$ accounts for the uncertainty in measurements as in [44],

$$\mathcal{R}_\epsilon(S_0(z)) := S_0(z) \cup \{d \in \mathcal{D} \mid \exists d' \in S_0(z), f(z, d') - Lq_z(d, d') - \epsilon \geq T_{\min} \\ \wedge f(z, d') + Lq_z(d, d') + \epsilon \leq T_{\max}\}. \quad (4)$$

We denote by \mathcal{R}_ϵ^n the n -time reachability operator, which calls \mathcal{R}_ϵ n times using the previous step's output. Then, $\lim_{n \rightarrow +\infty} \mathcal{R}_\epsilon^n(S_0(z))$ represents the subset of \mathcal{D} that can be identified as safe for the context z using the initial safe set $S_0(z)$, by observing f up to a statistical certainty restricted by ϵ .

4 Theoretical analyses

Consider a sequence of patient contexts $\bar{z} = [z_1 \dots z_N]$. Let $\mathbb{X}_N = X_1 \times \dots \times X_N$ denote the space of all context-admissible recommendation pairs, where $X_n = z_n \times \mathbb{D}_n$, and $\mathbb{D}_n \subseteq \mathcal{D}$ is the admissible dose space for z_n . For a given sequence of context-recommendation set A , let \mathbf{y}_A denote the $|A|$ -dimensional vector containing corresponding noisy evaluations of f . The quantity governing our regret bounds after N rounds in this scenario is a volatility-adapted maximum information gain term, $\gamma_N^{vol} = \max_{A \subseteq \mathbb{X}_N} I(\mathbf{y}_A; \mathbf{f}_A)$, where $\mathbf{f}_A = [f(\mathbf{x})]_{\mathbf{x} \in A}$ and $I(\mathbf{y}_A; \mathbf{f}_A)$ is the mutual information between f and observations at points in A . In the general setting where there is not a fixed context sequence, we have $\gamma_N = \max_{A \subseteq \mathcal{X}^N} I(\mathbf{y}_A; \mathbf{f}_A)$. Note that since $\mathbb{X}_N \subseteq \mathcal{X}^N$, we have $\gamma_N^{vol} \leq \gamma_N$. Explicit bounds on γ_N depending on N are studied in the literature [41, 51]. In this section, we first derive a high probability upper bound on the cumulative regret of TACO for a fixed context sequence without safety constraints. Then, we bound the regret of ESCADA in a single context scenario with safety constraints. For the former, we have $\bar{z}_1 = [z_1 \dots z_N]$, and $\mathbb{D}_n = \mathcal{D}$, and we denote the upper bound on the information gain term (see Lemma 2) by γ_N^{vol1} . For the latter, we have $\bar{z}_2 = [z \dots z]$, $\mathbb{D}_n = S_n(z)$, and we denote the upper bound on the information gain term by γ_N^{vol2} . We also prove that every dose recommended by ESCADA is safe with high probability (w.h.p.). Detailed proofs for each result can be found in Appendix D.

First, we mention two standard results. Lemma 1 shows that $f(x)$ is contained in the GP-induced confidence intervals w.h.p. and Lemma 2 expresses the information gain in terms of predictive variances.

Lemma 1. (Theorem 1 in [25]) Pick $\delta \in (0, 1)$, and define $\beta_n = 2L^2 + 300\gamma_n \log^3(n/\delta)$, where L is the Lipschitz constant. Let $\mathcal{E} = \{|\mu_{n-1}(\mathbf{x}) - f(\mathbf{x})| \leq \beta_n^{1/2} \sigma_{n-1}(\mathbf{x}), \forall n \in \mathbb{N}, \forall \mathbf{x} \in \mathcal{X}\}$. We have $\mathbb{P}\{\mathcal{E}\} \geq 1 - \delta$.

Lemma 2. (Lemma 5.3 in [41]) The information gain for the points selected can be expressed in terms of the predictive variances. If $\mathbf{f}_N = (f(\mathbf{x}_n))$, $I(\mathbf{y}_N; \mathbf{f}_N) = \frac{1}{2} \sum_{n=1}^N \log(1 + \sigma^{-2} \sigma_{n-1}^2(\mathbf{x}_n))$.

The following theorem provides a safety guarantee for ESCADA under the event \mathcal{E} in Lemma 1. The proof depends on an inductive argument on the safe sets constructed by ESCADA.

Theorem 1. Given that an initial safe dose set $S_0(\mathbf{z})$ is available $\forall \mathbf{z} \in \mathcal{Z}$, all doses recommended by ESCADA are safe, that is, $T_{\min} \leq f(\mathbf{z}_n, d_n) \leq T_{\max} \forall n \in [N]$, with at least $1 - \delta$ probability.

We proposed a novel acquisition function, TACO, for the *leveling* problem described in §2. Theorem 2 provides an upper bound on the regret of TACO without any safety constraints in place.

Theorem 2. Define β_n as in Lemma 1 and let $C := 8/\log(1 + \sigma^{-2})$. Cumulative regret of TACO for a fixed context sequence is upper-bounded as follows,

$$\mathbb{P}\{R_N \leq \sqrt{CN\beta_N\gamma_N^{vol1}}\} \geq 1 - \delta.$$

Next, we introduce a new concept, *safe path*.

Definition 1. (Safe Path) For a fixed context $\mathbf{z} \in \mathcal{Z}$, we say that there exists a safe path between two doses $d_1, d_2 \in \mathcal{D}$ if the following is satisfied,

$$\eta(d_1, d_2) = \min \left(\min_{d \in [d_1, d_2]} (T_{\max} - \epsilon - f(\mathbf{z}, d)), \min_{d \in [d_1, d_2]} (f(\mathbf{z}, d) - T_{\min} - \epsilon) \right) > 0, \quad (5)$$

where $\epsilon > 0$ is same as in (4). Definition 1 states that if there exists a safe path between two doses d_1 and d_2 , then there is no dose violating or exactly at the safety constraints between them. That is, $f(d) \in (T_{\min} + \epsilon + \eta(d_1, d_2), T_{\max} - \epsilon - \eta(d_1, d_2))$ for all $d \in [d_1, d_2]$. Next, we give the regret bound for ESCADA, which uses TACO with the safe sets $S_n(\mathbf{z}_n)$ in (3). We assume a fixed context scenario and show that the safety constraints result in at most a constant addition to the regret.

Theorem 3. If there exists a safe path between at least one dose $d \in S_0(\mathbf{z})$ and d_z^* , and we have $q_z(d_1, d_2) = K(|d_1 - d_2|)$ for some monotonically increasing mapping $K : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and for all $d_1, d_2 \in \mathcal{D}$, then the cumulative regret of ESCADA in a safety constrained single context (\mathbf{z}) scenario can be upper-bounded by setting the discretization parameter $\lambda < \epsilon$ as follows,

$$\mathbb{P}\{R_N \leq \sqrt{CN\beta_N\gamma_N^{vol2}} + \bar{T}N_z\} \geq 1 - \delta,$$

where $N_z \in \mathbb{N}$ is a constant independent of N .

Note that since $f(\mathbf{z}, d_z^*) = T$ and $T \in (T_{\min} + \alpha, T_{\max} - \alpha)$, one must ensure that $\alpha > \epsilon$ for the possibility of a safe path between some $d \in S_0(\mathbf{z})$ and d_z^* at the first place.

The assumption that $q_z(d_1, d_2) = K(|d_1 - d_2|)$ for a monotonically increasing mapping K holds in our working example where the blood glucose response to insulin dose can be characterized by the carbohydrate factor (CF) [38, 55]. That is, if we let $L \gg \text{CF}$, then we have $f(\mathbf{z}, d_1) - f(\mathbf{z}, d_2) \leq L|d_1 - d_2|$ for $d_1, d_2 \in \mathcal{D}$, and $q_z(d_1, d_2) = |d_1 - d_2|$. Moreover, this is the case for a variety of widely used kernel induced distance metrics. For the squared exponential kernel $k(\alpha, \beta) = \exp(-\|\alpha - \beta\|^2/2\sigma^2)$, we have (see §2),

$$q_z(d_1, d_2) = \sqrt{2 - 2\exp(-|d_1 - d_2|^2/\sigma^2)} \quad (6)$$

Similar observations follow for other radial-basis function kernels (e.g., Laplacian kernel). Theorems 2 and 3 constitute the non-incremental parts in our analysis as they provide explicit regret guarantees for a novel problem structure and acquisition strategy, both with and without safety constraints for a *compact* and *convex* action set. To generalize the bound in Theorem 3 to mixed context scenarios, one needs to impose further assumptions on the regularity of context arrivals over time. We provide experimental results on mixed context scenarios in §5 and show that the inter-contextual information transfer actually improves the performance as expected.

5 Experiments

5.1 Experimental setup

Online experimentation in the clinical setting is hazardous and it faces ethical challenges [12, 34, 35, 52]. Previous works on dose-finding clinical trials validate their methods either through synthetic experiments or by using external algorithms to fit a dose-response model to real-world data when

the patient group is homogeneous [5, 26, 40]. Such algorithms are not applicable in our case as they assume a shared dose-response model among patients, whereas we aim to learn *personalized* models. We make *in silico* experiments using the open-source implementation [57] of the U.S. FDA approved University of Virginia (UVA)/PADOVA T1DM simulator [23], which is the most frequently used framework in blood glucose control studies [10, 11, 13, 18, 28, 46, 59, 60]. It comes with 30 virtual patients with different individual characteristics: 10 adults, 10 adolescents, and 10 children. The simulator calculates the postprandial blood glucose (PPBG) response of a patient for (meal event, bolus insulin dose) pairs using differential equations and patient characteristics [23]. In our best effort to evaluate the success and potential of ESCADA as a supplementary tool in the clinical setting and to provide external validation, we also compare its performance against a clinician for five virtual adult patients. Our code is available at <https://github.com/Bilkent-CYBORG/ESCADA>.

Performance metrics When the PPBG level drops below 70 mg/dl (or exceeds 180 mg/dl), hypoglycemia (hyperglycemia) events occur. Both events may lead to life-threatening conditions [6]. Our primary objective is to recommend insulin doses that keep the patients’ PPBG level close to the target BG level (see (1)) while not recommending any insulin dose that triggers hypoglycemia or hyperglycemia events (see (2)). We set the target blood glucose (BG) level to 112.5 mg/dl [22]. We gauge an algorithm’s performance by combining its regret, hypoglycemia/hyperglycemia frequencies (error frequencies), and glycemic risk indices. Glycemic risk indices are low blood glycemic index (LBGI) and high blood glycemic index (HBGI), and they characterize the risk of hypoglycemia and hyperglycemia events in the long term, respectively [22]. A well-rounded algorithm should have a low cumulative regret together with small risk index values by *safely* and *efficiently* learning to recommend better insulin doses. Besides, we discuss the competing algorithms’ consistency since inexplicable variations in medical therapy are undesirable [47]. Precisely speaking, for a *fixed history*, when we query a recommendation from a consistent algorithm multiple times for the same meal event, it should not change. A meal event is a two-element tuple: (carbohydrate intake, fasting blood glucose). We create different meal events via uniform sampling to create an ensemble of different scenarios. We sample carbohydrate intake for each meal event from [20, 80] g, and fasting blood glucose from [100, 150] mg/dl.

Single meal event (SME) scenario In this part, we recommend insulin doses to a patient for the same meal event, assuming that the patient takes the insulin dose directly before the meal. Simulating this setup is helpful for two reasons: (i) it tests the performance of the algorithms in the classical *non-contextual* MAB setting, (ii) it provides a simple benchmark to understand the performance metrics and to compare them with the contextual setup later. Our objective is to optimize the PPBG 150 minutes after the meal. We make 15 consecutive dose recommendations for a meal event in a single run. We repeat this experiment with 30 different meal events for all 30 patients.

Multiple meal events (MME) scenario In this part, we recommend insulin doses to a single patient for a sequence of different meal events and use the same 30 meal events created in the SME scenario. We make consecutive recommendations for different meal events in a round-robin fashion and recommend a total of 15 doses for each meal event. Precisely speaking, after making a dose recommendation for a meal event, we make recommendations for the other 29 meal events and observe the PPBGs before making the next recommendation for the same meal event. This setup illustrates that the information gained from a context can assist in making decisions for different contexts. Contextual knowledge transfer enables our algorithm to adapt to intra- and inter-daily variability in meal events.

Algorithms We simulate ESCADA and TACO (i.e., without the safety mechanism). Besides, we propose a Thompson sampling (TS)-based algorithm and its safe version (STS), which operate as follows: TS samples a PPBG function from the posterior GP in each round and recommends the dose that achieves the PPBG closest to the target BG. STS implements the safe exploration strategy in §3 and uses TS as the acquisition function. In the final part, we implement the GP-UCB algorithm in [41] using three different “reward” functions in Figure 1 and compare it to our acquisition functions TACO and TS. We use two versions of dose calculators as baselines, whose details are given below.

Dose calculators Dose calculators are commonly used in diabetes care, as they are transparent and interpretable [55]. We use them to initialize the safe dose set for patient and meal event pairs. A calculator recommends an insulin dose via a simple equation, including carbohydrate intake, fasting blood glucose, and patient-specific parameters. They must be fine-tuned to ensure safety which may be challenging. Even when fine-tuned, they may not include some patient characteristics which can affect PPBG in the calculation rule. *Correction doses* constitute 9% of the patients’ daily insulin dose intake due to the calculator’s failure [55]. More details about bolus calculators are available in Appendix C.

Table 1: “-TC” indicates that tuned calculator was used. Target PPBG level is $T = 112.5$ mg/dl. “PPBG” column is averaged over observations for all 30 patients, 30 meal events per patient, and 15 recommendations per meal event. “Hyper” and “Hypo” columns denote the hyperglycemia and hypoglycemia event frequencies, respectively, averaged over all 30 patients. Similarly, HBGI and LBGI risk indices are averaged over all 30 patients. We report (mean \pm standard deviation).

	Algorithm	PPBG	Hyper	Hypo	HBGI	LBGI
	Calc.	144.0 ± 39.5	$.143 \pm .217$	$.0614 \pm .189$	3.84 ± 3.41	1.37 ± 3.95
	Tuned Calc.	123.7 ± 18.1	0	$.0021 \pm .010$	0.83 ± 0.66	0.24 ± 0.47
SME	TS	119.8 ± 42.2	$.046 \pm .029$	$.0216 \pm .028$	1.52 ± 1.25	1.01 ± 2.31
	TACO	121.7 ± 50.4	$.049 \pm .032$	$.0175 \pm .019$	1.89 ± 1.63	0.52 ± 0.46
	STS	121.6 ± 24.8	$.031 \pm .063$	$.0029 \pm .010$	1.07 ± 1.33	0.15 ± 0.23
	ESCADA	122.2 ± 20.0	$.015 \pm .030$	$.0031 \pm .008$	0.77 ± 0.82	0.11 ± 0.24
	STS-TC	117.1 ± 11.9	$.002 \pm .004$	$.0004 \pm .001$	0.28 ± 0.24	0.05 ± 0.05
	ESCADA-TC	116.1 ± 12.5	$.002 \pm .004$	$.0007 \pm .003$	0.26 ± 0.21	0.07 ± 0.09
MME	GP-UCB-1	124.1 ± 87.0	$.179 \pm .050$	$.2618 \pm .202$	5.43 ± 2.26	15.4 ± 22.5
	GP-UCB-2	103.7 ± 59.4	$.080 \pm .060$	$.2873 \pm .254$	2.21 ± 1.58	16.0 ± 27.0
	GP-UCB-3	111.0 ± 32.6	$.022 \pm .010$	$.0648 \pm .076$	0.73 ± 0.33	3.45 ± 5.17
	TS	112.4 ± 14.4	$.003 \pm .003$	$.0107 \pm .011$	0.16 ± 0.18	0.53 ± 0.95
	TACO	113.7 ± 19.2	$.006 \pm .031$	$.0010 \pm .002$	0.29 ± 1.18	0.07 ± 0.04
	STS	116.5 ± 12.5	$.004 \pm .015$	$.0007 \pm .002$	0.32 ± 0.49	0.05 ± 0.05
	ESCADA	116.9 ± 13.1	$.006 \pm .017$	$.0005 \pm .002$	0.34 ± 0.55	0.04 ± 0.04

We consider two setups. First, we use a calculator setting that occasionally fails to provide safe dose recommendations and sacrifice the assumption that an initial safe set, $S_0(z)$, is always available. Then, we use tuned calculators for each patient and ensure that $S_0(z)$ is *almost* always available.

5.2 Discussion of results

Safety Ensuring patient safety is pivotal. Theorem 1 shows that ESCADA recommends safe doses with high probability when an initial safe dose set is available. However, the initially provided set may not always be safe in reality due to calculator or clinician mistakes. We simulate two scenarios when an initial safe set is almost always available and not. For the latter, Table 1 shows that the error frequencies of ESCADA are not zero. We expect that error since the calculator fails to consistently provide safe doses in the beginning. However, ESCADA yields significantly lower error frequencies and risk index values than the calculator. That improvement stems from ESCADA’s ability to gradually identify and recommend safe doses, even when initially misdirected. We plot consecutive dose recommendations by ESCADA in SME scenario for three different meal events in Figure 4. For each of these meal events, rule-based calculator fails to provide safe doses in the beginning. Notwithstanding, ESCADA expands its safe set in the right direction and eventually recommends safe doses. Figure 3 and Table 1 confirm the safety mechanism’s effectiveness as ESCADA and STS yield significantly better safety metrics than the unsafe algorithms, TACO and TS, especially for hypoglycemia. Next, we manually tune the calculator parameters for each patient separately so that it successfully provides an initial safe set almost always (Tuned Calc., Table 1). Table 1 shows that ESCADA-TC and STS-TC yield remarkably lower error frequencies and risk indices, along with better PPBG distributions.

Regret Minimizing the regret is equivalent to recommending doses that lead to PPBG values close to the target BG by (1). We observe from Figures 3 and 5, and Table 1 that ESCADA(-TC) and STS(-TC) significantly outperforms the (tuned) calculator. Figure 5 shows that TACO and TS incur lower cumulative regrets than ESCADA and STS in the MME scenario. That is a natural trade-off between safety and regret since the safety mechanism restricts the allocation of a dose before it is identified as safe. Therefore, a safe algorithm yields higher regret when the initial safe set is far from the optimal dose.

Inter-contextual information transfer We investigate the efficiency of GP-induced smoothness in *transferring* information between different contexts. We mark an evident advancement in PPBG distributions and safety metrics in the MME scenario compared to the SME scenario in Table 1. Examining Figure 6, we observe that ESCADA expands the safe dose set and identifies the optimal dose faster in the MME scenario. Remember that ESCADA recommends doses for different meal




Figure 3: PPBG distribution boxplots in SME scenario. “-TC” suffix indicates that the tuned calculator is used.




Figure 4: Consecutive dose recommendations to three different meal events with *unsafe* $S_0(z)$ in SME scenario.




Figure 5: The cumulative regrets averaged over all 30 patients in MME scenario (± 0.25 standard deviation).




Figure 6: Consecutive dose recommendations for three different meal events (ME) in SME and MME scenarios.

events between two consecutive recommendations for the same meal event in the MME scenario. That is, the information gained from a context improves the performance for other contexts. Besides, we observe significant advances in the safety metrics of TACO and TS in the MME scenario as well.

Comparison with GP-UCB We compare our acquisition functions, TACO and TS, against the adaptations of the GP-UCB algorithm as described in §2 for three different reward functions in Figure 1. “GP-UCB-X” uses “Reward X” in Figure 1, which are defined as follows at each round n ,

$$r_1(n) = -\log(|y_n - T| + 1) \quad r_2(n) = 1 - \exp(-|y_n - T|/20) \quad r_3(n) = -|y_n - T|$$

We have y_n instead of $f(z_n, d_n)$ as the observations are noisy. Figures 7, 8, and 9 show that GP-UCB’s performance varies wildly for different rewards, and it is outperformed by TACO and TS. The practitioner needs to choose a “good” reward function for each problem. Our algorithms do not require that.

Consistency Figures 3 and 5, and Table 1 reveal that ESCADA and STS yield similar results. Both algorithms use GPs and have $\mathcal{O}(n^3)$ time and $\mathcal{O}(n^2)$ memory complexities where n is the number of observations. The key difference between them is that STS strikes the balance between exploration and exploitation through intrinsic randomization. That is, for a fixed patient history, STS can make different recommendations for the same meal event in test time, damaging its interpretability and leading to undesired inexplicable variations in the treatment [47]. On the other hand, ESCADA trades-off the exploration and exploitation through the explicit and deterministic machinery described in §3 which makes it a fairly interpretable model. Moreover, even though we design and test STS, we do not provide an upper bound on its regret as opposed to ESCADA, which is an interesting future work.

Clinician comparison We compare ESCADA’s performance against a clinician’s for five virtual patients. For each patient, we provided the clinician with 20 samples in the form of (meal event, insulin dose, PPBG) and asked her to make recommendations for 20 *unseen* meal events. We provided ESCADA with the same 20 samples for each patient and queried recommendations for the same 20 test meal events. Figure 10 shows that the clinician performs slightly worse than the calculator,

and ESCADA outperforms both significantly. These results suggest that making inferences about a patient's dose response is not trivial, and ESCADA is promising supplementary tool in clinical setting. Moreover, ESCADA can provide the clinicians with various useful statistics regarding dose responses, such as the confidence region of the response, hypo-/hyper-glycemia probabilities, or probability of response residing in a specific interval for a given patient, meal event, and insulin dose.




Figure 7: PPBG distribution boxplots for TACO, TS, GP-UCB, and the calculator in MME scenario.




Figure 8: Cumulative regrets for TACO, TS, GP-UCB, and calculator in MME scenario (± 0.25 standard deviation)




Figure 9: Cumulative regrets for TACO and the best GP-UCB in MME scenario (± 0.25 standard deviation).




Figure 10: PPBG distribution boxplots for five virtual patients in the clinician comparison experiment.

6 Concluding remarks

We formalized and studied a prevalent problem in medicine, *safe leveling*, and proposed TACO, a novel acquisition function tailored for this problem structure. As safety is crucial in healthcare, we proposed a safe exploration strategy to render TACO safe. Combining these two blocks, we proposed ESCADA, a *safe* and *efficient* learning algorithm, and provided safety guarantees and upper bounds on its cumulative regret. Through extensive *in silico* experiments on the bolus-insulin dose allocation problem for type-1 diabetes disease, we showed our algorithms' effectiveness over the rule-based dose calculators and straightforward adaptations of the GP-UCB algorithm for the *safe leveling* task. We also compared ESCADA's performance against a clinician's to provide external validation and discussed its potential as a complementary instrument in clinical settings. ESCADA can also be used in other safety-critical decision-making problems where the goal is to safely control a target variable.

Acknowledgments: This work was supported by the Scientific and Technological Research Council of Turkey (TUBITAK) under Grant 215E342. Ilker Demirel was also supported by Vodafone as part of 5G and Beyond Joint Graduate Support Programme coordinated by Information and Communication Technologies Authority. The clinician experiment was done as part of the TUBITAK Project 215E342 under the supervision of a Professor in the Dept. of Internal Diseases in Umraniye Training and Research Hospital, Istanbul, Turkey. Dose recommendations for the virtual patients are provided by a clinical dietician working in the same hospital. Research conducted within the scope of the TUBITAK Project 215E342 has been approved by the ethics committee of the hospital.

References

- [1] M. Abeille and A. Lazaric. Linear Thompson Sampling Revisited. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 176–184, 2017.
- [2] S. Amani, M. Alizadeh, and C. Thrampoulidis. Regret bound for safe Gaussian process bandit optimization. In *Conference on Learning for Dynamics and Control*, pages 158–159, 2020.
- [3] E. A. Ashley. Towards precision medicine. *Nature Reviews Genetics*, 17(9):507, 2016.
- [4] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2):235–256, 2002.
- [5] M. Aziz, E. Kaufmann, and M.-K. Riviere. On multi-armed bandit designs for dose-finding clinical trials. *Journal of Machine Learning Research*, 22:1–38, 2021.
- [6] A. Bastaki et al. Diabetes mellitus and its treatment. *International Journal of Diabetes and Metabolism*, 13(3):111, 2005.
- [7] M. Br  g  re, P. Gaillard, Y. Goude, and G. Stoltz. Target tracking for contextual bandits: Application to demand side management. In *International Conference on Machine Learning (ICML)*, pages 754–763, 2019.
- [8] A. R. Cardoso and H. Xu. Risk-averse stochastic convex bandit. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 39–47, 2019.
- [9] A. Cassel, S. Mannor, and A. Zeevi. A general approach to multi-armed bandits under risk criteria. In *Conference On Learning Theory (COLT)*, pages 1295–1306, 2018.
- [10] Y. Chandak, G. Theocharous, S. Shankar, M. White, S. Mahadevan, and P. Thomas. Optimizing for the future in non-stationary mdps. In *International Conference on Machine Learning (ICML)*, pages 1414–1425, 2020.
- [11] Y. Chandak, S. Niekum, B. da Silva, E. Learned-Miller, E. Brunskill, and P. S. Thomas. Universal off-policy evaluation. *Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- [12] I. Y. Chen, E. Pierson, S. Rose, S. Joshi, K. Ferryman, and M. Ghassemi. Ethical machine learning in healthcare. *Annual Review of Biomedical Data Science*, 4, 2020.
- [13] I. Fox, J. Lee, R. Pop-Busui, and J. Wiens. Deep reinforcement learning for closed-loop blood glucose control. In *Machine Learning for Healthcare Conference (ML4HC)*, pages 508–536, 2020.
- [14] N. Galichet, M. Sebag, and O. Teytaud. Exploration vs exploitation vs safety: Risk-aware multi-armed bandits. In *Asian Conference on Machine Learning*, pages 245–260, 2013.
- [15] M. A. Gelbart, J. Snoek, and R. P. Adams. Bayesian optimization with unknown constraints. In *Uncertainty in Artificial Intelligence (UAI)*, pages 250–259, 2014.
- [16] GPy. GPy: A Gaussian process framework in Python. url = <http://github.com/SheffieldML/GPy>, 2012.
- [17] J. M. Hernandez-Lobato, M. A. Gelbart, R. P. Adams, M. W. Hoffman, Z. Ghahramani, et al. A general framework for constrained Bayesian optimization using information-based search. *Journal of Machine Learning Research*, 17(160):1–53, 2016.
- [18] A. Huang, L. Leqi, Z. Lipton, and K. Azizzadenesheli. Off-policy risk assessment in contextual bandits. *Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- [19] A. Kazerouni, M. Ghavamzadeh, Y. Abbasi Yadkori, and B. Van Roy. Conservative contextual linear bandits. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- [20] N. Khajonchotpanya, Y. Xue, and N. Rujeerapaiboon. A revised approach for risk-averse multi-armed bandits under CVaR criterion. *Operations Research Letters*, 49(4):465–472, 2021.
- [21] K. Khezeli and E. Bitar. Safe linear stochastic bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10202–10209, 2020.
- [22] B. P. Kovatchev, D. J. Cox, A. Kumar, L. Gonder-Frederick, and W. L. Clarke. Algorithmic evaluation of metabolic control and risk of severe hypoglycemia in type 1 and type 2 diabetes using self-monitoring blood glucose data. *Diabetes Technology & Therapeutics*, 5(5):817–828, 2003.

- [23] B. P. Kovatchev, M. Breton, C. Dalla Man, and C. Cobelli. In silico preclinical trials: a proof of concept in closed-loop control of type 1 diabetes, 2009.
- [24] M. D. Kraft, I. F. Btaiche, G. S. Sacks, and K. A. Kudsk. Treatment of electrolyte disorders in adult patients in the intensive care unit. *American Journal of Health-System Pharmacy*, 62(16):1663–1682, 2005.
- [25] A. Krause and C. S. Ong. Contextual Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2447–2455, 2011.
- [26] H.-S. Lee, C. Shen, J. Jordon, and M. Schaar. Contextual constrained learning for dose-finding clinical trials. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 2645–2654, 2020.
- [27] H.-S. Lee, C. Shen, W. Zame, J.-W. Lee, and M. van der Schaar. Sdf-bayes: Cautious optimism in safe dose-finding clinical trials with drug combinations and heterogeneous patient groups. In *International Conference on Artificial Intelligence and Statistics AISTATS*, pages 2980–2988, 2021.
- [28] S. Lee, J. Kim, S. W. Park, S.-M. Jin, and S.-M. Park. Toward a fully automated artificial pancreas system using a bioinspired reinforcement learning design: In silico validation. *IEEE Journal of Biomedical and Health Informatics*, 25(2):536–546, 2020.
- [29] T. Lu, D. Pál, and M. Pál. Contextual multi-armed bandits. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 485–492, 2010.
- [30] R. Mirnezami, J. Nicholson, and A. Darzi. Preparing for precision medicine. *New England Journal of Medicine*, 366(6):489–491, 2012.
- [31] T. M. Moldovan and P. Abbeel. Safe exploration in Markov decision processes. In *International Conference on Machine Learning (ICML)*, 2012.
- [32] A. Moradipari, S. Amani, M. Alizadeh, and C. Thrampoulidis. Safe linear thompson sampling with side information. *IEEE Transactions on Signal Processing*, 2021.
- [33] K. A. Nerenberg, K. B. Zarnke, A. A. Leung, K. Dasgupta, S. Butalia, K. McBrien, K. C. Harris, M. Nakhla, L. Cloutier, M. Gelfer, et al. Hypertension Canada’s 2018 guidelines for diagnosis, risk assessment, prevention, and treatment of hypertension in adults and children. *Canadian Journal of Cardiology*, 34(5): 506–525, 2018.
- [34] W. Price and I. Nicholson. Regulating black-box medicine. *Mich. L. Rev.*, 116:421, 2017.
- [35] W. N. Price. Big data and black-box medical algorithms. *Science Translational Medicine*, 10(471), 2018.
- [36] C. E. Rasmussen. Gaussian processes in machine learning. *Advanced Lectures on Machine Learning*, 2004.
- [37] A. Sani, A. Lazaric, and R. Munos. Risk-aversion in multi-armed bandits. *Advances in Neural Information Processing Systems (NeurIPS)*, 25, 2012.
- [38] S. Schmidt and K. Nørgaard. Bolus calculators. *Journal of Diabetes Science and Technology*, 8(5): 1035–1041, 2014.
- [39] B. Schölkopf, A. J. Smola, F. Bach, et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT Press, 2002.
- [40] C. Shen, Z. Wang, S. Villar, and M. Van Der Schaar. Learning for dose allocation in adaptive clinical trials with safety constraints. In *International Conference on Machine Learning (ICML)*, pages 8730–8740, 2020.
- [41] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. In *International Conference on Machine Learning (ICML)*, pages 1015–1022, 2010.
- [42] M. C. Steinbach. Markowitz revisited: Mean-variance models in financial portfolio analysis. *SIAM Review*, 43(1):31–85, 2001.
- [43] I. Steinwart and A. Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- [44] Y. Sui, A. Gotovos, J. Burdick, and A. Krause. Safe exploration for optimization with Gaussian processes. In *International Conference on Machine Learning (ICML)*, pages 997–1005, 2015.

- [45] Y. Sui, Vincent Zhuang, J. Burdick, and Y. Yue. Stagewise safe Bayesian optimization with Gaussian processes. In *International Conference on Machine Learning (ICML)*, pages 4781–4789, 2018.
- [46] M. Tejedor, A. Z. Woldaregay, and F. Godtliebsen. Reinforcement learning application in diabetes blood glucose control: A systematic review. *Artificial Intelligence in Medicine*, 104:101836, 2020.
- [47] C. R. Tomson and S. N. Van Der Veer. Learning from practice variation to improve the quality of care. *Clinical Medicine*, 13(1):19, 2013.
- [48] M. Turchetta, F. Berkenkamp, and A. Krause. Safe exploration for interactive machine learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [49] S. Vakili and Q. Zhao. Mean-variance and value at risk in multi-armed bandit problems. In *Annual Allerton Conference on Communication, Control, and Computing*, pages 1330–1335. IEEE, 2015.
- [50] S. Vakili and Q. Zhao. Risk-averse multi-armed bandit problems under mean-variance measure. *IEEE Journal of Selected Topics in Signal Processing*, 10(6):1093–1111, 2016.
- [51] S. Vakili, K. Khezeli, and V. Picheny. On information gain and regret bounds in gaussian process bandits. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 82–90, 2021.
- [52] E. Vayena, A. Blasimme, and I. G. Cohen. Machine learning in medicine: addressing ethical challenges. *PLoS Medicine*, 15(11):e1002689, 2018.
- [53] S. S. Villar and W. F. Rosenberger. Covariate-adjusted response-adaptive randomization for multi-arm clinical trials using a modified forward looking Gittins index rule. *Biometrics*, 74(1):49–57, 2018.
- [54] A. Wachi, Y. Sui, Y. Yue, and M. Ono. Safe exploration and optimization of constrained MDPs using Gaussian processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [55] J. Walsh, R. Roberts, and T. Bailey. Guidelines for optimal bolus calculator settings in adults. *Journal of Diabetes Science and Technology*, 5(1):129–135, 2011.
- [56] Y. Wu, R. Shariff, T. Lattimore, and C. Szepesvári. Conservative bandits. In *International Conference on Machine Learning (ICML)*, pages 1254–1262, 2016.
- [57] J. Xie. Simglucose v0.2.1. url=<https://github.com/jxx123/simglucose>, 2018.
- [58] Y. Zhang, E. B. Laber, M. Davidian, and A. A. Tsiatis. Interpretable dynamic treatment regimes. *Journal of the American Statistical Association*, 113(524):1541–1549, 2018.
- [59] T. Zhu, K. Li, P. Herrero, and P. Georgiou. Basal glucose control in type 1 diabetes using deep reinforcement learning: An in silico validation. *IEEE Journal of Biomedical and Health Informatics*, 2020.
- [60] T. Zhu, K. Li, L. Kuang, P. Herrero, and P. Georgiou. An insulin bolus advisor for type 1 diabetes using deep reinforcement learning. *Sensors*, 20(18):5058, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) See the discussions after (6), and the comments on Safe Thompson Sampling (STS) in **Consistency** part in §5.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#) See §2.
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) See Appendix D.
3. If you ran experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) Complete code repository is included in the supplementary material with a README file containing the instructions on reproduction.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) See §5 and the code repository (includes comments).
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) See Table 1 and Figures 3, 5, 7, 8, and 9.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) See Appendix E.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#) See §5 and Appendix E. We cite [16] and [57].
 - (b) Did you mention the license of the assets? [\[Yes\]](#) See Appendix E for the licences of the used assets.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[Yes\]](#) We share our complete code repository for the experiments in the supplementary material, which will be made publicly available after publishing the manuscript.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)

A Related work

Table 2: Comparison of ESCADA to some other safe exploration algorithms for GP optimization.

Safety mechanism	Safety during exploration	Efficiency*	Two-sided safety constraint	New acquisition with regret bounds
[2]	✓	✓	✗	✗
[15]	✗	✓	✗	✗
[17]	✗	✓	✓	✗
[45]	✓	✗	✗	✗
[44]	✓	✗	✗	✗
[48]	✓	✓	✗	✗
ESCADA (Ours)	✓	✓	✓	✓


*The algorithm does not designate the safe set exploration as a proxy objective. That is, it does not make queries whose sole purpose is to expand the safe set.

Table 2 highlights the important aspects of the discussion in Section 1 on related works and their comparison against ESCADA. ESCADA tackles a novel problem structure defined in Section 2 by employing a novel acquisition strategy, TACO, together with an efficient safe exploration scheme for a two-sided safety constraint.


B Additional Experimental Results

Figure 11a and Table 3 present the results for the single meal event (SME) scenario, separately for three different subgroups which are “Adult”, “Adolescent”, and “Child”. The target blood glucose (BG) level is set to $T = 112.5$ mg/dl for all the subgroups for this experiment. While the resulting PPBG distributions are similar among the subgroups, the safety constraints are better satisfied for adults and adolescents than for children.

In Figure 11b, we present the bolus-insulin dose recommendation distributions in the SME scenario, separately for the “Adult”, “Adolescent”, and “Child” subgroups. While the safety bounds and the target blood glucose level are the same for all the subgroups, we observe significantly different dose recommendation distributions. Such an observation is not surprising since each individual needs a different treatment strategy to achieve the same objective. Together with Figure 11a, Figure 11b demonstrates our algorithms’ ability to provide personalized treatment strategies.



(a) PPBG distributions.




(b) Bolus-insulin dose recommendation distributions.


Figure 11: PPBG and bolus-insulin dose recommendation distributions for different subgroups in the SME scenario. Target BG level is the same for all the subgroups and it is $T = 112.5$ mg/dl.

In practice, the clinician may want to set different objectives for different patients. Our algorithms are flexible in the sense that the clinician can easily set the safety bounds T_{\min} and T_{\max} , and the target level T separately for each patient, without having to tune any further parameters. Note that, however, our objective is not to find some optimal T_{\min} , T_{\max} , T values, but to learn to provide dose recommendations according to the values determined by the clinician.

To further corroborate our algorithms' ability for providing personalized treatment strategies, we set different target BG levels for every patient subgroup. Note that we can choose different targets for every individual as well. However, for the sake of presentation, we limit the difference to be between the subgroups only. We set $T = 140$ mg/dl for the adults, $T = 110$ mg/dl for the adolescents, and $T = 125$ mg/dl for the children. Inspecting Figure 12a and Table 4, we see that our algorithms can adapt to different target BG levels successfully. The difference between dose recommendation distributions in Figures 11b and 12b also shows that the algorithms can learn different strategies when the target BG level T is changed, for instance, by the clinician in practice.



(a) PPBG distributions.



(b) Bolus-insulin dose recommendation distributions.

Figure 12: PPBG and bolus-insulin dose recommendation distributions for different subgroups in the SME scenario. Target BG level is $T = 140$ mg/dl for the adults, $T = 110$ mg/dl for the adolescents, and $T = 125$ mg/dl for the children.

Table 3: SME scenario with the same target BG level for all the subgroups, $T = 112.5$ mg/dl.

	Algorithm	PPBG	Hyper	Hypo	HBGI	LBGI
Adult	Calc.	155.9 \pm 19.8	.158 \pm .201	0	4.09 \pm 2.19	0
	Tuned Calc.	134.7 \pm 12.0	0	0	1.36 \pm 0.62	0.002 \pm 0.007
	TS	117.4 \pm 26.4	.035 \pm .017	.012 \pm .020	0.85 \pm 0.42	0.53 \pm 0.68
	TACO	117.0 \pm 29.0	.041 \pm .027	.019 \pm .020	0.92 \pm 0.62	0.53 \pm 0.42
	STS	124.9 \pm 20.6	.022 \pm .036	0	1.07 \pm 0.78	0.05 \pm 0.04
	ESCADA	123.7 \pm 15.4	.007 \pm .012	0	0.72 \pm 0.46	0.02 \pm 0.02
	STS-TC	118.2 \pm 8.7	0	0	0.24 \pm 0.07	0.01 \pm 0.01
	ESCADA-TC	117.7 \pm 9.3	0	0	0.24 \pm 0.01	0.03 \pm 0.01
Adolescent	Calc.	153.9 \pm 33.8	.132 \pm .265	0	4.51 \pm 4.38	0.08 \pm 0.24
	Tuned Calc.	122.7 \pm 16.9	0	0	0.73 \pm 0.53	0.17 \pm 0.27
	TS	119.6 \pm 33.0	.036 \pm .029	.014 \pm .016	1.16 \pm 1.00	0.49 \pm 0.50
	TACO	122.1 \pm 45.1	.049 \pm .043	.016 \pm .017	1.78 \pm 1.80	0.50 \pm 0.48
	STS	125.1 \pm 26.8	.044 \pm .099	0	1.39 \pm 2.02	0.07 \pm 0.05
	ESCADA	124.8 \pm 19.0	.019 \pm .047	0	0.93 \pm 1.22	0.02 \pm 0.03
	STS-TC	116.6 \pm 10.2	.0002 \pm .0007	0	0.23 \pm 0.28	0.05 \pm 0.03
	ESCADA-TC	115.9 \pm 10.5	0	0	0.22 \pm 0.25	0.07 \pm 0.03
Child	Calc.	122.3 \pm 49.3	.140 \pm .177	.184 \pm .291	2.94 \pm 3.10	4.04 \pm 6.00
	Tuned Calc.	113.8 \pm 18.4	0	.006 \pm .017	0.41 \pm 0.41	0.55 \pm 0.65
	TS	122.5 \pm 59.6	.067 \pm .026	.038 \pm .035	2.55 \pm 1.38	2.06 \pm 3.70
	TACO	126.0 \pm 68.5	.059 \pm .014	.018 \pm .021	2.97 \pm 1.50	0.54 \pm 0.49
	STS	114.9 \pm 25.2	.028 \pm .026	.009 \pm .016	0.75 \pm 0.61	0.34 \pm 0.32
	ESCADA	117.9 \pm 24.0	.019 \pm .016	.009 \pm .013	0.66 \pm 0.53	0.31 \pm 0.33
	STS-TC	116.5 \pm 15.5	.006 \pm .007	.001 \pm .002	0.37 \pm 0.28	0.10 \pm 0.06
	ESCADA-TC	114.7 \pm 16.4	.006 \pm .006	.002 \pm .005	0.31 \pm 0.23	0.14 \pm 0.13

Table 4: SME scenario with the target level $T = 140$ mg/dl for the adults, $T = 110$ mg/dl for the adolescents, and $T = 125$ mg/dl for the children.

	Algorithm	PPBG	Hyper	Hypo	HBGI	LBGI
Adult	Calc.	176.7 \pm 19.8	.380 \pm .304	0	7.36 \pm 2.55	0
	Tuned Calc.	146.3 \pm 12.4	0	0	0.83 \pm 0.66	0.24 \pm 0.47
	TS	143.6 \pm 29.1	.066 \pm .056	.015 \pm .022	2.60 \pm 1.00	0
	TACO	145.6 \pm 32.5	.081 \pm .119	.013 \pm .012	3.18 \pm 1.62	0.53 \pm 0.74
	STS	145.4 \pm 19.3	.071 \pm .107	0	2.67 \pm 1.51	0
	ESCADA	148.2 \pm 16.0	.054 \pm .099	0	2.91 \pm 1.06	0
	STS-TC	139.4 \pm 8.3	0	0	1.70 \pm 0.17	0
	ESCADA-TC	140.7 \pm 8.4	0	0	1.84 \pm 0.16	0
Adolescent	Calc.	152.2 \pm 33.8	.120 \pm .261	0	4.30 \pm 4.31	0.10 \pm 0.29
	Tuned Calc.	121.9 \pm 16.9	.0002 \pm .0007	0	0.69 \pm 0.55	0.18 \pm 0.29
	TS	116.6 \pm 34.9	.036 \pm .026	.020 \pm .023	1.15 \pm 0.91	0.67 \pm 0.62
	TACO	119.3 \pm 46.6	.051 \pm .046	.023 \pm .019	1.80 \pm 1.91	0.66 \pm 0.57
	STS	123.9 \pm 27.0	.043 \pm .098	0	1.36 \pm 1.95	0.11 \pm 0.07
	ESCADA	123.0 \pm 19.4	.018 \pm .047	0	0.87 \pm 1.21	0.03 \pm 0.04
	STS-TC	115.1 \pm 10.7	0	0	0.22 \pm 0.34	0.06 \pm 0.03
	ESCADA-TC	114.4 \pm 11.0	0	0	0.21 \pm 0.29	0.09 \pm 0.03
Child	Calc.	135.6 \pm 52.2	.272 \pm .291	.101 \pm .176	4.37 \pm 4.44	2.26 \pm 3.40
	Tuned Calc.	119.8 \pm 17.0	0	0	0.60 \pm 0.38	0.20 \pm 0.20
	TS	133.0 \pm 55.5	.067 \pm .036	.025 \pm .015	2.83 \pm 1.52	0.72 \pm 0.43
	TACO	134.1 \pm 56.1	.047 \pm .026	.010 \pm .012	2.61 \pm 1.70	0.28 \pm 0.28
	STS	123.8 \pm 18.0	.023 \pm .032	.005 \pm .009	0.76 \pm 0.59	0.13 \pm 0.16
	ESCADA	127.4 \pm 15.5	.016 \pm .018	.009 \pm .012	0.85 \pm 0.32	0.20 \pm 0.24
	STS-TC	126.5 \pm 12.3	.006 \pm .006	.0006 \pm .0014	0.70 \pm 0.12	0.04 \pm 0.03
	ESCADA-TC	126.2 \pm 12.7	.006 \pm .009	.001 \pm .004	0.69 \pm 0.17	0.04 \pm 0.05

C Rule-based bolus insulin dose calculators

Rule-based bolus insulin dose calculators are commonly used in diabetes care, as they are transparent and interpretable [55]. We use them to determine the initial safe dose (i.e., $S_0(\mathbf{z})$) for each patient. A rule based-calculator recommends a bolus insulin dose via a simple equation. We use the one from [38], which is given below.

$$\text{Bolus-insulin Dose} = \left(\frac{\text{CHO}}{\text{ICR}} + \frac{G_M - G_T}{\text{CF}} \right)^+ \quad (7)$$

where $a^+ := \max\{0, a\}$, CHO (g) is the carbohydrate intake, ICR (g/U) is the insulin-to-carbohydrate ratio, and CF (mg/dl/U) is the insulin correction factor. G_M and G_T (mg/dl) denote the preprandial BG level and the postprandial target BG levels, respectively. In order to prevent hypoglycemia and hyperglycemia events, ICR and CF values need to be precisely tuned for each patient, which may prove to be challenging in reality. Even if ICR and CF values are tuned correctly, rule-based calculators (e.g., (7)) discount other patient characteristics which may affect PPBG. Research shows that the correction doses constitute 9% of the patients' total daily insulin dose intake due to the calculator's failure [55].

D Proofs

D.1 Proof of Theorem 1

We prove by induction when f is L -Lipschitz continuous and the event \mathcal{E} in Lemma 1 holds. When \mathcal{E} holds, we have $f(\mathbf{z}_n, d) \in [\mu_{n-1}(\mathbf{z}_n, d) - \beta_n^{1/2} \sigma_{n-1}(\mathbf{z}_n, d), \mu_{n-1}(\mathbf{z}_n, d) + \beta_n^{1/2} \sigma_{n-1}(\mathbf{z}_n, d)]$, that is,

$$l_n(\mathbf{z}_n, d) \leq f(\mathbf{z}_n, d) \leq u_n(\mathbf{z}_n, d),$$

for every $d \in \mathcal{D}$ and $\mathbf{z}_n \in \mathcal{Z}$. Also, by Lipschitz continuity, we have,

$$\begin{aligned} f(\mathbf{z}_n, d) &\geq f(\mathbf{z}_n, d') - Lq_{\mathbf{z}_n}(d, d') \\ &\geq l_n(\mathbf{z}_n, d') - Lq_{\mathbf{z}_n}(d, d'), \end{aligned}$$

and,

$$\begin{aligned} f(\mathbf{z}_n, d) &\leq f(\mathbf{z}_n, d') + Lq_{\mathbf{z}_n}(d, d') \\ &\leq u_n(\mathbf{z}_n, d') + Lq_{\mathbf{z}_n}(d, d'), \end{aligned}$$

that is,

$$l_n(\mathbf{z}_n, d') - Lq_{\mathbf{z}_n}(d, d') \leq f(\mathbf{z}_n, d) \leq u_n(\mathbf{z}_n, d') + Lq_{\mathbf{z}_n}(d, d'),$$

for any $d, d' \in \mathcal{D}$.

Remember, we have,

$$\bar{l}_n(\mathbf{z}, d) = \max \{l_n(\mathbf{z}, d), l_n(\mathbf{z}, d') - Lq_{\mathbf{z}}(d, d')\} \quad (8)$$

$$\bar{u}_n(\mathbf{z}, d) = \min \{u_n(\mathbf{z}, d), u_n(\mathbf{z}, d') + Lq_{\mathbf{z}}(d, d')\}, \quad (9)$$

where $d' = \operatorname{argmin}_{d^* \in \overline{\mathcal{D}}_{\mathbf{z}}} q_{\mathbf{z}}(d, d^*)$. Then, we have, $f(\mathbf{z}_n, d) \geq \bar{l}_n(\mathbf{z}_n, d)$ and $f(\mathbf{z}_n, d) \leq \bar{u}_n(\mathbf{z}_n, d)$. For the base case, $n = 0$, we have $f(\mathbf{z}, d) \in [T_{\min}, T_{\max}]$ for any $d \in S_0(\mathbf{z})$ by definition. For the inductive step, first consider the safety condition from below, that is, $f(\mathbf{z}_n, d_n) \geq T_{\min}$. For $n \geq 1$, assume that $f(\mathbf{z}_n, d) \geq T_{\min}$ for any $d \in S_{n-1}(\mathbf{z}_n)$. Then by definition of the safe set expansion operator, for all $d^* \in S_n(\mathbf{z}_n) \setminus S_{n-1}(\mathbf{z}_n)$, there exists $d' \in S_{n-1}(\mathbf{z}_n)$ such that,

$$\begin{aligned} T_{\min} &\leq \bar{l}_n(\mathbf{z}_n, d') - Lq_{\mathbf{z}_n}(d', d^*) \\ &\leq f(\mathbf{z}_n, d') - Lq_{\mathbf{z}_n}(d', d^*) \\ &\leq f(\mathbf{z}_n, d^*). \end{aligned} \quad (10)$$

Similarly, to check the condition from above, we can observe,

$$\begin{aligned} T_{\max} &\geq \bar{u}_n(\mathbf{z}_n, d') + Lq_{\mathbf{z}_n}(d', d^*) \\ &\geq f(\mathbf{z}_n, d') + Lq_{\mathbf{z}_n}(d', d^*) \\ &\geq f(\mathbf{z}_n, d^*), \end{aligned} \quad (11)$$

where (10) and (11) follow from Lipschitz continuity, and we are done.

D.2 Proof of Theorem 2

Assume \mathcal{E} in Lemma 1 holds. Then,

$$\begin{aligned} R_N &= \sum_{n=1}^N |f(\mathbf{z}_n, d_n) - T| \\ &\leq \sum_{n=1}^N (|f(\mathbf{z}_n, d_n) - \mu_{n-1}(\mathbf{z}_n, d_n)| + |(\mu_{n-1}(\mathbf{z}_n, d_n) - T)|) \end{aligned} \quad (12)$$

$$\leq \sum_{n=1}^N 2\beta_n^{1/2}\sigma_{n-1}(\mathbf{z}_n, d_n), \quad (13)$$

where (12) follows from triangle inequality on \mathbb{R} . For the first term in (12), we have,

$$|f(\mathbf{z}_n, d_n) - \mu_{n-1}(\mathbf{z}_n, d_n)| \leq \beta_n^{1/2}\sigma_{n-1}(\mathbf{z}_n, d_n),$$

by Lemma 1 under the good event \mathcal{E} . For the second term, notice that in the first step of TACO, we choose the doses whose confidence intervals contain the target value T , i.e., $\bar{l}_n(\mathbf{z}_n, d_n) \leq T \leq \bar{u}_n(\mathbf{z}_n, d_n)$. By (8) and (9), we have,

$$l_n(\mathbf{z}_n, d_n) \leq T \leq u_n(\mathbf{z}_n, d_n),$$

which is equivalent to,

$$\mu_{n-1}(\mathbf{z}_n, d_n) - \beta_n^{1/2}\sigma_{n-1}(\mathbf{z}_n, d_n) \leq T \leq \mu_{n-1}(\mathbf{z}_n, d_n) + \beta_n^{1/2}\sigma_{n-1}(\mathbf{z}_n, d_n), \quad (14)$$

by definitions of $l_n(\mathbf{z}_n, d_n)$ and $u_n(\mathbf{z}_n, d_n)$ in Section 3. Finally, we have, $|\mu_{n-1}(\mathbf{z}_n, d_n) - T| \leq \beta_n^{1/2}\sigma_{n-1}(\mathbf{z}_n, d_n)$ by (14). Note that since there is not any safety constraint, we have at least one dose in $\mathbb{D}_n = \mathcal{D}$ (e.g., $d_{\mathbf{z}_n}^*$) whose confidence interval contains the target value, T , which makes the passage from (12) to (13) always possible. Then, by Cauchy-Schwartz inequality,

$$\begin{aligned} R_N^2 &\leq N \sum_{n=1}^N 4\beta_n\sigma_{n-1}^2(\mathbf{z}_n, d_n) \\ &\leq 4N\beta_N\sigma^2 \sum_{n=1}^N \sigma^{-2}\sigma_{n-1}^2(\mathbf{z}_n, d_n) \end{aligned} \quad (15)$$

$$\leq 8N\beta_N\sigma^2 \frac{1}{2} \sum_{n=1}^N \frac{\sigma^{-2}\log(1 + \sigma^{-2}\sigma_{n-1}^2(\mathbf{z}_n, d_n))}{\log(1 + \sigma^{-2})} \quad (16)$$

$$\begin{aligned} &\leq CN\beta_N \max \frac{1}{2} \sum_{n=1}^N \log(1 + \sigma^{-2}\sigma_{n-1}^2(\mathbf{z}_n, d_n)) \\ &\leq CN\beta_N\gamma_N^{vol1}, \end{aligned} \quad (17)$$

where (15) holds since β_n is increasing in n . (16) follows from the fact that $s^2 < \frac{\sigma^{-2}\log(1+s^2)}{\log(1+\sigma^{-2})}$ for $s \in [0, \sigma^{-2}]$, and $\sigma^{-2}\sigma_{n-1}^2(\mathbf{z}_n, d_n) \leq \sigma^{-2}k(\mathbf{x}_n, \mathbf{x}_n) \leq \sigma^{-2}$, and (17) follows from combining Lemma 2 and the definition of γ_N^{vol1} in Section 4. Theorem 2 follows from taking square roots of both sides.

D.3 Auxiliary Results

Before we prove Theorem 3, it is useful to highlight the critical difference between the setup considered in Theorem 2 and Theorem 3, and motivate the proof of Theorem 3. In Theorem 2, note that the passage between (12) and (13) is always possible. This is because it is always guaranteed that there exists at least one dose $d \in \mathbb{D}_n$, whose confidence interval contains the target, T , where $\mathbb{D}_n = \mathcal{D}$ (e.g. $d_{\mathbf{z}_n}^*$). That allows us to upper-bound the term $|\mu_{n-1}(\mathbf{z}_n, d_n) - T|$ by $\beta_n^{1/2}\sigma_{n-1}(\mathbf{z}_n, d_n)$. However, for the setup considered in Theorem 3, we are not allowed to choose any $d \in \mathcal{D}$, but choose $d \in S_n(\mathbf{z}_n)$ in round n , to guarantee the safety of the recommendation. Until we are certain that the

optimal dose $d_z^* \in S_n(z)$, we will incur a linear regret at the worst case. Once we have $d_z^* \in S_n(z)$, the regret will converge to the same order as in Theorem 1, since the relation between (12) and (13) will be valid.

In the next three lemmas, we show that if $d_z^* \in \mathcal{R}_\epsilon^m(S_0(z))$ for some $m \in \mathbb{N}$ (i.e., d_z^* is reachable from the initial safe set $S_0(z)$ after finitely many iterations), we have $d_z^* \in S_n(z)$ after a finite number of rounds, and then the passage between (12) and (13) is also possible for the safety constrained setting in a single context scenario.

Lemma 3. *Posterior variance of the Gaussian process after round n at $\mathbf{x} \in \mathcal{X}$ can be upper bounded by the noise variance and the number of times the function evaluated at \mathbf{x} up to round n ($n_{\mathbf{x}}$) as, $\sigma^2/n_{\mathbf{x}} \geq \sigma_{n-1}^2(\mathbf{x})$.*

Proof. Given $A \subseteq B$, $H(\mu|A) \geq H(\mu|B)$, since conditioning on more observations will reduce entropy. Let $\mathbf{Y}_n^{\mathbf{x}}$ denote the observations made at $\mathbf{x} \in \mathcal{X}$ up to round n , and \mathbf{Y}_n denote all the observations up to round n . Then, we have $\mathbf{Y}_n^{\mathbf{x}} \subseteq \mathbf{Y}_n$, which implies $H(\mu|\mathbf{Y}_n^{\mathbf{x}}) \geq H(\mu|\mathbf{Y}_n)$. The entropy of a Gaussian random variable is $H(\mathcal{N}(\mu, \sigma^2)) = \frac{1}{2} \log(2\pi e \sigma^2)$. We have $H(\mu|\mathbf{Y}_n) = \frac{1}{2} \log(2\pi e \sigma_{n-1}^2(\mathbf{x}))$, and $H(\mu|\mathbf{Y}_n^{\mathbf{x}}) = \frac{1}{2} \log(2\pi e (n_{\mathbf{x}}/\sigma^2 + \sigma_0^{-2})^{-1})$, where $\sigma_0^2 = k(\mathbf{x}, \mathbf{x})$ is the prior variance of GP at \mathbf{x} . Then we can write,

$$\frac{1}{2} \log\left(\frac{2\pi e}{n_{\mathbf{x}}/\sigma^2 + \sigma_0^{-2}}\right) \geq \frac{1}{2} \log(2\pi e \sigma_{n-1}^2(\mathbf{x})) . \quad (18)$$

Since $k(\mathbf{x}, \mathbf{x}) \geq 0$, the following is immediate from (18),

$$\frac{\sigma^2}{n_{\mathbf{x}}} \geq \sigma_{n-1}^2(\mathbf{x}) .$$

□

Let us first to recall the definition of a *safe path*.

Definition 1. (Safe Path) For a fixed context $\mathbf{z} \in \mathcal{Z}$, we say that there exists a safe path between two doses $d_1, d_2 \in \mathcal{D}$ if the following is satisfied,

$$\eta(d_1, d_2) = \min \left(\min_{d \in [d_1, d_2]} (T_{\max} - \epsilon - f(\mathbf{z}, d)), \min_{d \in [d_1, d_2]} (f(\mathbf{z}, d) - T_{\min} - \epsilon) \right) > 0 , \quad (19)$$

Lemma 4. *For a fixed context $\mathbf{z} \in \mathcal{Z}$, if there exists at least one dose $d \in S_0(\mathbf{z})$ such that there exists a safe path between d and d_z^* , and we have $q_{\mathbf{z}}(d^1, d^2) = K(|d^1 - d^2|)$ for some monotonically increasing mapping $K : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and for all $d^1, d^2 \in \mathcal{D}$, then we have $d_z^* \in \mathcal{R}_\epsilon^m(S_0(\mathbf{z}))$ for some $m \in \mathbb{N}$.*

Proof. We first note that $\mathcal{R}_\epsilon^l(A) \subseteq \mathcal{R}_\epsilon^l(B)$ for any $l \in \mathbb{N}$ if $A \subseteq B$ by definition of the reachability operator in (4). That is, $\mathcal{R}_\epsilon^k(\{d\}) \subseteq \mathcal{R}_\epsilon^k(S_0(\mathbf{z}))$ since $d \in S_0(\mathbf{z})$. At the heart of the proof lies the idea that if there is a safe path between d and d_z^* , the reachability operator will keep expanding towards d_z^* . Now, if $d = d_z^*$, we are already done. Consider the case where $d_z^* > d$. Let $d_1 > d$ be such that,

$$Lq_{\mathbf{z}}(d, d_1) = \eta , \quad (20)$$

where $\eta := \eta(d, d_z^*)$. Then, by (19), we have,

$$\begin{aligned} Lq_{\mathbf{z}}(d, d_1) &\leq T_{\max} - \epsilon - f(\mathbf{z}, d) \\ \Rightarrow f(\mathbf{z}, d) + Lq_{\mathbf{z}}(d, d_1) + \epsilon &\leq T_{\max} , \end{aligned} \quad (21)$$

and,

$$\begin{aligned} Lq_{\mathbf{z}}(d, d_1) &\leq f(\mathbf{z}, d) - T_{\min} - \epsilon \\ \Rightarrow f(\mathbf{z}, d) - Lq_{\mathbf{z}}(d, d_1) - \epsilon &\geq T_{\min} . \end{aligned} \quad (22)$$

By (21), (22), and (4), we have $d_1 \in \mathcal{R}_\epsilon^1(\{d\})$, where $d_1 \geq d$ and $q_{\mathbf{z}}(d, d_1) = \eta/L$. Now, if $q_{\mathbf{z}}(d, d_1) = K(|d - d_1|)$ for some monotonically increasing mapping $K : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, where

$K(x) = 0 \iff x = 0$ (since $q_{\mathbf{z}}(\cdot, \cdot)$ is a metric), then we have $|d - d_1| = K^{-1}(\eta/L)$, which implies that,

$$d_1 = d + K^{-1}(\eta/L), \quad (23)$$

and $d_1 > d$ since $\eta > 0$. Now, if $d_1 \geq d_{\mathbf{z}}^* > d$, one can simply check that $Lq_{\mathbf{z}}(d, d_{\mathbf{z}}^*) \leq \eta$ by (20) since K is monotonically increasing. This then implies that $d_{\mathbf{z}}^* \in \mathcal{R}_{\epsilon}^1(\{d\})$, and we are done. If $d_{\mathbf{z}}^* > d_1$, we will keep expanding the reachable set. Note that we have $\mathcal{R}_{\epsilon}^1(\{d_1\}) \subseteq \mathcal{R}_{\epsilon}^2(\{d\})$, since $\{d_1\} \subseteq \mathcal{R}_{\epsilon}^1(\{d\})$, where $\mathcal{R}_{\epsilon}^2(\{d\}) = \mathcal{R}_{\epsilon}^1(\mathcal{R}_{\epsilon}^1(\{d\}))$. Similar to before, we can show that there exists $d_2 \in \mathcal{D}$ such that $d_2 \in \mathcal{R}_{\epsilon}^1(\{d_1\})$ and,

$$\begin{aligned} d_2 &= d_1 + K^{-1}(\eta/L) \\ &= d + 2K^{-1}(\eta/L), \end{aligned} \quad (24)$$

which follows from (23). By this iteration, we have $d_n \in \mathcal{R}_{\epsilon}^n(\{d\})$, where

$$d_n = d + nK^{-1}(\eta/L), \quad (25)$$

for $n \in \mathbb{N}$ such that $d_n \leq d_{\mathbf{z}}^*$. Let $M := |d_{\mathbf{z}}^* - d|$, and $m \in \mathbb{N}$ be the smallest integer such that $mK^{-1}(\eta/L) \geq M$. This implies that $d_{m-1} < d_{\mathbf{z}}^*$. Then, by (25), after one more iteration, we have $d_m \geq d_{\mathbf{z}}^*$, and $d_{\mathbf{z}}^* \in \mathcal{R}_{\epsilon}^m(\{d\})$, which completes the proof for the case where $d_{\mathbf{z}}^* > d$. For $d_{\mathbf{z}}^* < d$, similar arguments will follow. \square

Consider the fixed context scenario where the contexts are $\mathbf{z} \in \mathcal{Z}$ for all rounds $n \in \{1, \dots, N\}$. We define the bad event in round n as

$$\mathcal{G}_n = \{T \notin [\bar{l}_n(\mathbf{z}, d), \bar{u}_n(\mathbf{z}, d)], \forall d \in S_n(\mathbf{z})\}. \quad (26)$$

When the event \mathcal{G}_n occurs in round n , we incur the following instantaneous worst-case regret,

$$|f(\mathbf{z}, d_n) - T| \leq \bar{T}, \quad (27)$$

since $f \in [0, \bar{T}]$. When the event $\neg \mathcal{G}_n$ occurs, the instantenous regret can be bounded similar to before when the event \mathcal{E} holds, that is,

$$\begin{aligned} |f(\mathbf{z}, d_n) - T| &\leq |f(\mathbf{z}, d_n) - \mu_{n-1}(\mathbf{z}, d_n)| + |\mu_{n-1}(\mathbf{z}, d_n) - T| \\ &\leq 2\beta_n^{1/2}\sigma_{n-1}(\mathbf{z}, d_n), \end{aligned} \quad (28)$$

The next lemma shows that the number of rounds in which \mathcal{G}_n can occur in a single context scenario is bounded (independent of time horizon N) given that the optimal dose $d_{\mathbf{z}}^*$ is ϵ -reachable from $S_0(\mathbf{z})$ after a finite number of rounds.

Lemma 5. *Assume that the event \mathcal{E} in Lemma 1 holds. Then, if $d_{\mathbf{z}}^* \in \mathcal{R}_{\epsilon}^k(S_0(\mathbf{z}))$ for some $k \in \mathbb{N}$, the event \mathcal{G}_n , $n \geq 1$, can only occur finitely many times.*

Proof. If we can show that after finite number of rounds that the bad event \mathcal{G}_n occurs, we have $d_{\mathbf{z}}^* \in S_n(\mathbf{z})$, we are done since it is guaranteed that there exists at least one dose in $S_n(\mathbf{z})$ ($d_{\mathbf{z}}^*$) whose confidence interval contains the target value. Assume to the contrary that the bad event \mathcal{G}_n occurs infinitely many times. We first show that there exists $n_1 \in \mathbb{N}$ such that $\mathcal{R}_{\epsilon}^1(S_0(\mathbf{z})) =: \mathcal{R}_{\epsilon}^1 \subseteq S_{n_1}(\mathbf{z})$. Assume to the contrary. When \mathcal{G}_n occurs, TACO chooses a dose from the finite set $\bar{S}_n(\mathbf{z}) := S_n(\mathbf{z}) \cap \bar{\mathcal{D}}_{\mathbf{z}} \subseteq \mathcal{R}_{\epsilon}^1 \cap \bar{\mathcal{D}}_{\mathbf{z}}$ in round n . In that scenario, let $\bar{S}(\mathbf{z})$ denote the biggest finite set such that $\bar{S}_n(\mathbf{z}) \subseteq \bar{S}(\mathbf{z}) \subseteq \mathcal{R}_{\epsilon}^1 \cap \bar{\mathcal{D}}_{\mathbf{z}}$ for every $n \in \mathbb{N}$, and let $\bar{S}_{m_1}(\mathbf{z}) = \bar{S}(\mathbf{z})$. TACO chooses the dose with the maximum confidence interval in $\bar{S}_n(\mathbf{z})$ when \mathcal{G}_n occurs. Then, this implies that the posterior variances of all doses in $\bar{S}_{m_1}(\mathbf{z})$ will go to zero by Lemma 3. For any $d \in S_n(\mathbf{z})$, the lower and upper confidence bounds are calculated as follows,

$$\begin{aligned} \bar{l}_n(\mathbf{z}, d) &= \max\{l_n(\mathbf{z}, d), l_n(\mathbf{z}, d') - Lq_{\mathbf{z}}(d, d')\} \\ \bar{u}_n(\mathbf{z}, d) &= \min\{u_n(\mathbf{z}, d), u_n(\mathbf{z}, d') + Lq_{\mathbf{z}}(d, d')\}, \end{aligned}$$

where $d' = \operatorname{argmin}_{d^* \in \bar{\mathcal{D}}_{\mathbf{z}}} q_{\mathbf{z}}(d, d^*)$. Then, we have,

$$\begin{aligned} w_n(\mathbf{z}, d) &\leq u_n(\mathbf{z}, d') + Lq_{\mathbf{z}}(d, d') - (l_n(\mathbf{z}, d') - Lq_{\mathbf{z}}(d, d')) \\ &= 2\beta_n^{1/2}\sigma_{n-1}(\mathbf{z}, d') + 2Lq_{\mathbf{z}}(d, d'), \end{aligned} \quad (29)$$

where $\sigma_{n-1}(\mathbf{z}, d')$ goes to zero for every $d' \in \bar{S}_{m_1}(\mathbf{z})$. Also, by definition of $\bar{\mathcal{D}}_{\mathbf{z}}$, we have $2Lq_{\mathbf{z}}(d, d') < 2L\frac{\lambda}{2L} = \lambda$. Finally, by choosing $\lambda < \epsilon$ and by (29), we can conclude that there exists $n_1 - 1 > m_1$ such that $w_n(\mathbf{z}, d) < \lambda < \epsilon$ for every $d \in S_n(\mathbf{z})$ for $n \geq n_1 - 1$. That is, we have, $f(\mathbf{z}, d) - \epsilon < \bar{l}_{n_1-1}(\mathbf{z}, d)$ and $f(\mathbf{z}, d) + \epsilon > \bar{u}_{n_1-1}(\mathbf{z}, d)$, for every $d \in S_{n_1-1}(\mathbf{z})$, where $S_0(\mathbf{z}) \subseteq S_{n_1-1}(\mathbf{z})$.

Now, let us restate the safe set expansion rule in (3) and the reachability operator in (4),

$$S_n(\mathbf{z}_n) = S_{n-1}(\mathbf{z}_n) \cup \left(\bigcup_{d \in S_{n-1}(\mathbf{z}_n)} \{d' \in \mathcal{D} \mid \bar{l}_n(\mathbf{z}_n, d) - Lq_{\mathbf{z}_n}(d, d') \geq T_{\min} \right. \\ \left. \wedge \bar{u}_n(\mathbf{z}_n, d) + Lq_{\mathbf{z}_n}(d, d') \leq T_{\max} \} \right),$$

and,

$$\mathcal{R}_\epsilon(S_0(\mathbf{z})) := S_0(\mathbf{z}) \cup \{d \in \mathcal{D} \mid \exists d' \in S_0(\mathbf{z}), f(\mathbf{z}, d') - Lq_{\mathbf{z}}(d, d') - \epsilon \geq T_{\min} \\ \wedge f(\mathbf{z}, d') + Lq_{\mathbf{z}}(d, d') + \epsilon \leq T_{\max}\}.$$

For every dose $d \in \mathcal{R}_\epsilon^1 \cap S_0(\mathbf{z})$, we have $d \in S_{n_1}(\mathbf{z})$ since $S_0(\mathbf{z}) \subseteq S_{n_1}(\mathbf{z})$. Now, for every dose $d \in \mathcal{R}_\epsilon^1 \setminus S_0(\mathbf{z})$, there exists $d' \in S_0(\mathbf{z})$ such that $f(\mathbf{z}, d') - Lq_{\mathbf{z}}(d, d') - \epsilon \geq T_{\min}$, and $f(\mathbf{z}, d') + Lq_{\mathbf{z}}(d, d') + \epsilon \leq T_{\max}$. Then, since $\bar{l}_{n_1-1}(\mathbf{z}, d') - Lq_{\mathbf{z}}(d, d') \geq f(\mathbf{z}, d') - Lq_{\mathbf{z}}(d, d') - \epsilon \geq T_{\min}$, and $\bar{u}_{n_1-1}(\mathbf{z}, d') + Lq_{\mathbf{z}}(d, d') \leq f(\mathbf{z}, d') + Lq_{\mathbf{z}}(d, d') + \epsilon \leq T_{\max}$, we can conclude that for every $d \in \mathcal{R}_\epsilon^1$, we have $d \in S_{n_1}(\mathbf{z})$, that is, $\mathcal{R}_\epsilon^1 \subseteq S_{n_1}(\mathbf{z})$ for some $n_1 \in \mathbb{N}$, which is a contradiction. After that point, we can define $\mathcal{R}_\epsilon^2 = \mathcal{R}_\epsilon(\mathcal{R}_\epsilon^1)$, where $\mathcal{R}_\epsilon^1 \subseteq S_{n_1}(\mathbf{z})$, and show that there exists $n_2 \geq n_1$ such that $\mathcal{R}_\epsilon^2 \subseteq S_{n_2}(\mathbf{z})$ and so on. Finally, we have $N_z = n_k$ such that $\mathcal{R}_\epsilon^k \subseteq S_{N_z}$, where $N_z \in \mathbb{N}$, meaning that $d_z^* \in S_n(\mathbf{z})$ for every $n \geq N_z$, which completes the proof (Note that $\mathcal{R}_\epsilon^k := \mathcal{R}_\epsilon^k(S_0(\mathbf{z}))$). \square

D.4 Proof of Theorem 3

For an event \mathcal{G} , let $\mathbb{I}\{\mathcal{G}\} = 1$ if \mathcal{G} holds and 0 otherwise. Assume that the event \mathcal{E} in Lemma 1 holds,

$$R_N = \sum_{n=1}^N |f(\mathbf{z}, d_n) - T| \\ = \sum_{n=1}^N \mathbb{I}\{\mathcal{G}\} \times |f(\mathbf{z}, d_n) - T| + \sum_{n=1}^N \mathbb{I}\{\neg\mathcal{G}\} \times |f(\mathbf{z}, d_n) - T| \\ \leq \sum_{n=1}^N \mathbb{I}\{\neg\mathcal{G}\} \times |f(\mathbf{z}, d_n) - T| + \bar{T}N_z \quad (30)$$

$$\leq \sum_{n=1}^N \mathbb{I}\{\neg\mathcal{G}\} \times (|f(\mathbf{z}, d_n) - \mu_{n-1}(\mathbf{z}, d_n)| + |(\mu_{n-1}(\mathbf{z}, d_n) - T)|) + \bar{T}N_z \quad (31)$$

$$\leq \sum_{n=1}^N \mathbb{I}\{\neg\mathcal{G}\} \times 2\beta_n^{1/2}\sigma_{n-1}(\mathbf{z}, d_n) + \bar{T}N_z, \quad (32)$$

where (30) follows from (27) and Lemma 5, (31) follows from triangle inequality on \mathbb{R} , and (32) follows from (28). The rest of the proof is similar to Theorem 2.

Let $\bar{R}_N := \sum_{n=1}^N \mathbb{I}\{\neg \mathcal{G}\} \times 2\beta_n^{1/2} \sigma_{n-1}(\mathbf{z}, d_n)$. Then, by Cauchy-Schwartz inequality,

$$\begin{aligned}
\bar{R}_N^2 &\leq N \sum_{n=1}^N \mathbb{I}\{\neg \mathcal{G}\} \times 4\beta_n \sigma_{n-1}^2(\mathbf{z}, d_n) \\
&\leq 4N\beta_N \sigma^2 \sum_{n=1}^N \sigma^{-2} \sigma_{n-1}^2(\mathbf{z}, d_n) \\
&\leq 8N\beta_N \sigma^2 \frac{1}{2} \sum_{n=1}^N \frac{\sigma^{-2} \log(1 + \sigma^{-2} \sigma_{n-1}^2(\mathbf{z}, d_n))}{\log(1 + \sigma^{-2})} \\
&\leq CN\beta_N \max \frac{1}{2} \sum_{n=1}^N \log(1 + \sigma^{-2} \sigma_{n-1}^2(\mathbf{z}, d_n)) \\
&= CN\beta_N \gamma_N^{vol2}, \tag{33}
\end{aligned}$$

where (33) follows from combining Lemma 2 and the definition of γ_N^{vol2} in Section 4. Taking square roots of both sides, we have $\bar{R}_N \leq \sqrt{CN\beta_N \gamma_N^{vol2}}$. Finally, combining (32) with the definition of \bar{R}_N , we have $R_N \leq \sqrt{CN\beta_N \gamma_N^{vol2}} + \bar{T}N_z$.

E Training details

Our training session includes calculating posterior distributions for Gaussian processes for different models (i.e., models in the rows of Table 1) and sampling functions from the posterior distributions. We use GPy library in Python programming language for GP implementations which is licensed under the BSD 3-Clause License [16]. We also used [57] for the open source Python implementation of UVa/PADOVA T1DM simulator, which is licensed under MIT License. We did not use any GPUs and trained all the models using a system with AMD Ryzen 5 3600x @3.8-4.4GHz AM4 CPU and 16GB RAM. The total amount of training to reproduce the results provided in the paper takes under 4 hours.