



HAL
open science

Counterfactual Learning of Continuous Stochastic Policies

Houssam Zenati, Alberto Bietti, Matthieu Martin, Eustache Diemert, Julien Mairal

► **To cite this version:**

Houssam Zenati, Alberto Bietti, Matthieu Martin, Eustache Diemert, Julien Mairal. Counterfactual Learning of Continuous Stochastic Policies. 2020. hal-02883423v1

HAL Id: hal-02883423

<https://hal.science/hal-02883423v1>

Preprint submitted on 29 Jun 2020 (v1), last revised 19 Aug 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Counterfactual Learning of Continuous Stochastic Policies

Houssam Zenati^{1,2}, Alberto Bietti³, Matthieu Martin¹, Eustache Diemert¹, and
Julien Mairal²

¹Criteo AI Lab

²Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

³Inria, DI, Ecole normale supérieure (PSL Research University), Paris, France

Abstract

Counterfactual reasoning from logged data has become increasingly important for many applications such as web advertising or healthcare. In this paper, we address the problem of counterfactual risk minimization (CRM) for learning a stochastic policy with continuous actions, whereas most existing work has focused on the discrete setting. Switching from discrete to continuous action spaces presents several difficulties as naive discretization strategies have been shown to perform poorly. To deal with this issue, we first introduce an effective contextual modelling strategy that learns a joint representation of contexts and actions based on positive definite kernels. Second, we empirically show that the optimization perspective of CRM is more important than previously thought, and we demonstrate the benefits of proximal point algorithms and differentiable estimators. Finally, we propose an evaluation protocol for offline policies in real-world logged systems, which is challenging since policies cannot be replayed on test data, and we release a new large-scale dataset along with multiple synthetic, yet realistic, evaluation setups.¹

1 Introduction

Logged interaction data is widely available in a variety of real-life problems such as drug dosage prescription [16], evaluation of recommender systems [19], or setting reserve prices in online auctions [6]. In these settings, an important task is to leverage past data in order to find a good *policy* for selecting actions (e.g., drug doses or bids) from available features (or *contexts*), rather than relying on randomized trials or sequential exploration, which may be costly or even sometimes unethical.

More precisely, we consider offline logged bandit feedback data, consisting of contexts, actions selected by a given *logging policy*, associated to observed rewards. This is known as *bandit feedback*, since the reward is only observed for the action chosen by the logging policy. The problem of finding a good policy thus requires a form of *counterfactual* reasoning to estimate what the rewards would have been, had we used a different policy. While this setting is non-sequential, we assume that learning a *stochastic* policy is required so as to later deploy it, and possibly gather new data. Indeed, when the logging policy is stochastic, one may obtain unbiased estimates of the expected reward under a new policy through importance sampling with the inverse propensity scoring (IPS) method [12]. Even though the variance of the IPS estimator is potentially unbounded, one may nevertheless use such a framework to optimize new policies without the need for costly experiments [6, 8, 29, 30].

In this paper, we focus on stochastic policies with continuous actions, which have received little attention in the context of counterfactual policy optimization [7, 16]. As noted in [16], the continuous case presents new difficulties in comparison to the discrete setting, and naive discretization appears to perform poorly,

¹The code with open-source implementations for experimental reproducibility is available at <https://github.com/criteo-research/optimization-continuous-action-crm>.

which was confirmed by our experiments. Our first contribution is then to introduce an effective contextual modelling using a joint nonlinear embedding of continuous actions and contexts, by using a finite-dimensional approximate embedding of a positive definite kernel [32].

Our second contribution is to underline the role of optimization algorithms for solving non-convex counterfactual risk minimization (CRM) problems [6, 30]. We believe that this aspect was overlooked, as previous work has mostly studied the effectiveness of *estimation* methods regardless of the optimization procedure. In this paper, we show that appropriate tools can bring significant benefits in practice. To that effect, we introduce differentiable estimators based on soft-clipping the importance weights, which are more amenable to gradient-based optimization than hard clipping procedures that have been used previously [6, 31]. We also find that the use of proximal point algorithms [10, 23, 25], which consist of approximately solving a sequence of regularized subproblems for minimizing a non-convex objective, tend to dominate simpler off-the-shelf optimization approaches.

Finally, we propose an offline evaluation protocol for logged bandit feedback data, when we cannot replay policies at test time; this is a challenging but useful setup, which, to the best of our knowledge, has never been formally addressed in the literature. This protocol relies on (i) a self-normalized importance sampling estimator, (ii) importance sampling diagnostics to discard unreliable solutions; and (iii) significance tests to assess improvements to a reference policy. We evaluate the quality of the protocol on synthetic datasets, and release a new real-life logged dataset of continuous actions with more than 100 millions of samples. This dataset will allow the research community to perform counterfactual policy evaluation with real-life data and can be found in the Appendix.

2 Related Work

A large effort in counterfactual risk minimization has been devoted to designing estimators that have less variance than the IPS method, through clipping importance weights [6, 31], variance regularization [29], or by leveraging reward estimators through doubly robust (DR) methods [8, 24].

In order to tackle an overfitting phenomenon termed “propensity overfitting”, Swaminathan and Joachims [30] also consider self-normalized estimators [22]. Such estimation techniques also appear in the context of sequential learning in contextual bandits, where the goal of the agent is to find a good policy in a sequential manner while minimizing regret [2, 17], as well as for off-policy evaluation in reinforcement learning [13]. In contrast, the setting considered in our work is not sequential.

While most approaches for counterfactual policy optimization tend to focus on discrete actions, few works have tackled the continuous action case, again with a focus on estimation rather than optimization. In particular, propensity scores for continuous actions were considered first in [11]. More recently, evaluation and optimization of continuous action policies were studied in a non-parametric context in [16] using kernel smoothing, and in [7] in a semi-parametric setting.

Note that these methods consider deterministic policies while our focus is on stochastic ones, but the kernel smoothing approach of [16] can be interpreted as learning stochastic policies with a noise level depending on the kernel bandwidth [9]—an analogy which we leverage in our work.

Optimization methods for learning stochastic policies have been mainly studied in the context of reinforcement learning through the policy gradient theorem [3, 28, 33]. Such methods typically need to observe samples from the new policy at each optimization step, which is not allowed in our off-policy setting. Other methods leverage a form of off-policy estimates during optimization for improving the policy at each step [15, 26], but this still requires fresh samples, while we consider objective functions involving only a fixed dataset of collected data. In the context of CRM, [27] introduces an estimator with a continuous clipping objective that achieves an improved bias-variance trade-off over DR. Nevertheless, this estimator is non-smooth, unlike our soft-clipping estimator.

3 Counterfactual Learning with Continuous Actions

In this section we begin by reviewing the *counterfactual risk minimization* framework, and then present our new modelling approaches for stochastic policies with continuous actions.

3.1 Background

For a stochastic policy $\pi : \mathcal{X} \rightarrow \Delta(\mathcal{A})$ over a set of actions \mathcal{A} , a contextual bandit environment generates i.i.d. context features $x \sim \mathcal{D}_X$ in \mathcal{X} , actions $a \sim \pi(\cdot|x)$ and feedbacks/losses $y \sim D_Y(\cdot|x, a)$ with expectation $\mathbb{E}[y|x, a] = \eta^*(x, a)$. We denote the resulting distribution over triples (x, a, y) as \mathcal{D}_π . We consider a logged dataset $(x_i, a_i, y_i, \pi_{0,i})$, $i = 1, \dots, n$, where we assume $(x_i, a_i, y_i) \sim \mathcal{D}_{\pi_0}$ i.i.d. for a given stochastic logging policy π_0 , and the propensities are denoted by $\pi_{0,i} := \pi_0(a_i|x_i)$. The expected loss or risk of a policy π is then given by

$$R(\pi) = \mathbb{E}_{(x,a,y) \sim \mathcal{D}_\pi} [y]. \quad (1)$$

For the logged bandit, the task is to determine a policy $\hat{\pi}$ in a set of *stochastic* policies Π with small risk. Note that this definition may also include deterministic policies by allowing Dirac masses (in fact, the minimizer of (1), given by $\pi^*(x) = \arg \min_a \eta^*(x, a)$, is deterministic), but it may be desirable to enforce a minimum variance in order to perform future offline experiments, assuming $\hat{\pi}$ is meant to be deployed. In our setting, the expectation in (1) cannot be computed directly for any π , as the available interaction data comes from a different distribution π_0 . Yet, multiple empirical estimators of the risk hereafter allow to derive an empirical optimal policy that is found by solving:

$$\hat{\pi} \in \arg \min_{\pi \in \Pi} \hat{R}(\pi) + \Omega(\pi), \quad (2)$$

where the objective consists of an empirical estimate of the risk and of possible data-dependent regularizers on the policy, denoted by Ω . When using counterfactual estimators for \hat{R} , this method has been called (regularized) *counterfactual risk minimization* [29].

The counterfactual approach tackles the mismatch between the logging policy $\pi_0(\cdot|x)$ and a policy π in Π via importance sampling. Assuming π_0 has non-zero mass on the support of π , the IPS method [12] relies on correcting the distribution mismatch to derive an unbiased empirical estimate

$$R(\pi) = \mathbb{E}_{(x,a,y) \sim \mathcal{D}_{\pi_0}} \left[y \frac{\pi(a|x)}{\pi_0(a|x)} \right], \quad \hat{R}_{\text{IPS}}(\pi) = \frac{1}{n} \sum_{i=1}^n y_i \frac{\pi(a_i|x_i)}{\pi_{0,i}}. \quad (3)$$

Since the empirical estimator $\hat{R}_{\text{IPS}}(\pi)$ in Eq. (3) has large variance and is subject to various overfitting phenomena [30], different regularization strategies have been proposed, such as clipping the importance sampling (cIPS) weights [6], or the *self-normalized* (SNIPS) estimator [30], see Appendix 8 for a review of these estimators and technical details.

Another approach is to directly fit the loss values in observed data with a direct estimator $\hat{\eta}(x, a)$, for instance by using ridge regression to fit $y_i \approx \hat{\eta}(x_i, a_i)$, and then to use the deterministic greedy policy $\hat{\pi}(x) = \arg \min_a \hat{\eta}(x, a)$. This approach, termed direct method (DM), has the benefit of avoiding the high-variance problems of IPS-based methods, but may suffer from large bias since it ignores the potential mismatch between $\hat{\pi}$ and π_0 .

Nevertheless, such loss estimators can be effective when few samples are available, and may be combined with IPS estimators (see Appendix 8) in the so-called doubly robust (DR) estimator [8].

3.2 Continuous Action Policy Parameterization

When considering continuous action spaces, the choice of policies is more involved than in the discrete action case, since the outcome of the policy may take infinitely many possible values, instead of a finite number, and the metric structure of the actions also needs to be taken into account. While one may leverage discrete

action strategies by first discretizing the action space, this may be suboptimal as the metric information of the action space is lost, and the policies obtained by predicting a finite set of actions will be discontinuous. Besides, the logged propensities $\pi_{0,i}$ need to be tailored to the specific discretization in order to have unbiased estimators, which may be impractical.

In this paper, we thus focus on stochastic policies belonging to certain classes of continuous distributions, such as normal or log-normal. Specifically, we consider a set of continuous policies $\Pi = \{\pi_\theta, \theta \in \Theta\}$ parameterized by $\theta = (\theta_\mu, \sigma)$, with mean $\mu_{\theta_\mu}(x)$ and standard deviation σ , for instance $\pi_\theta(\cdot|x) = \mathcal{N}(\mu_{\theta_\mu}(x), \sigma^2)$. As a simple baseline, one may take $\mu_{\theta_\mu}(x) = \theta_\mu^\top \varphi(x)$, where φ could be constant, piecewise constant, linear or non-linear. While such baselines may be effective in some simple problems, they may be limited in more difficult scenarios where the expected loss $\eta^*(x, a)$ has a complex behavior as a function of a , which motivates the need for classes of policies which can better capture such variability by considering x and a jointly.

The counterfactual loss predictor (CLP). One approach that tries to better model the behavior of $\eta^*(x, a)$ is the direct method, which directly estimates a loss predictor $\hat{\eta} \approx \eta^*$ using regression, leading to a policy $\pi(x) = \arg \min_a \hat{\eta}(x, a)$, but suffers from large bias since it is not trained in a counterfactual manner. Alternatively, we may consider (stochastic) policies of a similar form, with mean $\mu_\eta(x) = \arg \min_a \eta(x, a)$, and learn η using counterfactual risk minimization. Unfortunately, such policies are impractical since they rely on a difficult optimization problem on actions, and this in turn makes the corresponding CRM problem non-differentiable. Instead, we introduce a softmax approximation of μ_η using anchor points a_1, \dots, a_m , which we call *counterfactual loss predictor*, that avoids the need for optimization over actions, and makes the CRM problem differentiable:

$$\mu_{\text{CLP}, \eta}^\gamma(x) = \sum_{i=1}^m a_i \frac{\exp(-\gamma \eta(x, a_i))}{\sum_{j=1}^m \exp(-\gamma \eta(x, a_j))}. \quad (4)$$

Loss predictors based on kernels. The above CLP model is parameterized by regressors $\eta(x, a)$ that are meant to approximate the true loss $\eta^*(x, a)$. In a continuous action problem, a reasonable assumption is that $\eta^*(x, a)$ is smooth as a function of a . Then, a good choice is to take η in a space of smooth functions such as the reproducing kernel Hilbert space defined by a positive definite kernel, so that one may control the smoothness of η through regularization with the RKHS norm. More precisely, we consider kernels of the form $K((x, a), (x', a')) = K_{\mathcal{X}}(x, x')K_{\mathcal{A}}(a, a')$, which correspond to a tensor product feature map $\varphi(x, a) = \varphi_{\mathcal{X}}(x) \otimes \varphi_{\mathcal{A}}(a)$. For simplicity, we consider linear or quadratic kernels on contexts, leading to $\varphi_{\mathcal{X}}(x) = x$ or $\varphi_{\mathcal{X}}(x) = (xx^T, x)$, noting that more complex kernels may also be used. For actions, however, it is important to use a richer kernel such as the Gaussian kernel, which allows us to be fully non-parametric and approximate arbitrary feedback functions. For computational efficiency, we rely on a Nyström approximation of $\varphi_{\mathcal{A}}$ instead of using the exact kernel, which amounts to replacing $\varphi_{\mathcal{A}}$ by its projection on a finite-dimensional subspace of the RKHS $\mathcal{F} = \text{span}\{(\varphi_{\mathcal{A}}(a_1), \dots, \varphi_{\mathcal{A}}(a_m))\}$ for a set of anchor points $\bar{A} = \{a_1, \dots, a_m\}$. In practice we may choose a_i to be quantiles of observed actions, and can take them equal to the a_i in (4). The Nyström approximation then provides us a finite-dimensional embedding $\psi_{\mathcal{A}}(a) = K_{AA}^{-1/2} K_A(a)$, where $K_{AA} = [K_{\mathcal{A}}(a_i, a_j)]_{ij}$ and $K_A(a) = [K_{\mathcal{A}}(a, a_i)]_i$, see [32]. Then, learning functions of the form $\eta(x, a) = \langle \theta_\mu, \varphi_{\mathcal{X}}(x) \otimes \psi_{\mathcal{A}}(a) \rangle$ with appropriate ℓ^2 regularization corresponds to controlling the RKHS norm, and hence the smoothness of η . In particular, even with a very small number of anchor points, we may obtain a good approximation of the full kernel method, particularly when η^* is smooth, while other strategies based on discretizing the action space may require a much finer discretization, and tend to be less robust to this choice, as verified in our experiments (see Section 6).

4 Optimization-Driven Approaches for CRM

We now present optimization aspects of the CRM problem, as well as a new differentiable estimator.

Soft Clipping IPS. The hard clipping estimator $\hat{R}_{\text{cIPS}}^M(\pi) = \frac{1}{n} \sum_{i=1}^n y_i \min\{\pi(a_i|x_i)/\pi_{0,i}, M\}$ makes the objective function non-differentiable, and also causes terms in the objective with clipped weights to have zero gradient. In other words, a trivial stationary point of the objective function is that of a stochastic policy that differs enough from the logging policy such that all importance weights are clipped. To that effect, we propose a differentiable logarithmic soft-clipping strategy, which is helpful to escape such bad stationary points. Given a threshold parameter $M \geq 0$ and an importance weight $w_i = \pi(a_i|x_i)/\pi_{0,i}$, we consider the soft-clipped weights:

$$\zeta(w_i, M) = \begin{cases} w_i & \text{if } w_i \leq M \\ \alpha(M) \log(w_i + \alpha(M) - M) & \text{otherwise,} \end{cases} \quad (5)$$

where $\alpha(M)$ is such that $\alpha(M) \log(\alpha(M)) = M$, which yields a differentiable operator. We give further explanations about the benefits of clipping strategies as well as an illustration of this soft logarithmic clipping in Appendix 9.2. Then, the IPS estimator with soft clipping becomes

$$\hat{R}_{\text{scIPS}}^M(\pi) = \frac{1}{n} \sum_{i=1}^n y_i \zeta\left(\frac{\pi(a_i|x_i)}{\pi_{0,i}}, M\right). \quad (6)$$

In Appendix 13, we prove that the variance-regularized version of this estimator enjoys a similar generalization bound to that of Swaminathan and Joachims [29] for the hard-clipped version, and hence provides a good optimization objective for minimizing the expected risk.

Proposition 4.1 (Generalization bound for \hat{R}_{scIPS}^M). *Assume losses y_i in $[-1, 0]$ and importance weights bounded by W . With probability $1 - \delta$, we have, for all π in Π ,*

$$R(\pi) \leq \hat{R}_{\text{scIPS}}^M(\pi) + O\left(\sqrt{\frac{\hat{V}(\pi)C_n(\Pi, \delta)}{n}} + \frac{SC_n(\Pi, \delta)}{n}\right),$$

where $S = \zeta(W, M) = O(\log W)$, \hat{V} denotes the empirical variance of the loss estimates, and $C_n(\Pi, \delta)$ is a measure of complexity of the policy class (see Appendix 13).

Note that the result requires bounded importance weights in contrast to hard-clipping, but the dependence on the upper bound W is very mild (logarithmic), and we gain significant benefits in terms of optimization by having a smooth objective.

Proximal Point Algorithms. The non-convexity of the CRM objective (2) may come from various reasons such as the weight clipping strategies, variance regularization [30], the use of self-normalized estimators [30], or simply the policy parametrization. The resulting problems have then been optimized with classical gradient-based methods [29, 30] such as L-BFGS [20], or the stochastic gradient descent approach [14, 29].

Proximal point methods are classical approaches originally designed for convex optimization [25], which were then found to be useful for the minimization of nonconvex functions [10, 23]. In order to minimize a function \mathcal{L} , the main idea is to approximately solve a sequence of subproblems that are better conditioned than \mathcal{L} , such that the sequence of iterates converges towards a stationary point of the original problem. More precisely, the proximal point method consists of computing a sequence

$$\theta_k \approx \arg \min_{\theta} \left(\mathcal{L}(\theta) + \frac{\kappa}{2} \|\theta - \theta_{k-1}\|_2^2 \right), \quad (7)$$

where $\mathcal{L}(\theta) = \hat{R}(\pi_\theta) + \Omega(\pi_\theta)$ and $\kappa > 0$ is a constant parameter. By regularizing the loss with a quadratic function, the subproblems become intuitively “less nonconvex”. For many machine learning formulations, it is even possible to obtain convex sub-problems, as long as κ is large enough (see [23]). Therefore, instead of applying a gradient-based algorithm directly on \mathcal{L} , we will also consider the proximal point strategy (7) with a parameter κ , which we set to zero only for the last iteration.

Note that the effect of the proximal point algorithm differs from the proximal policy optimization (PPO) strategy used in reinforcement learning [26], even though both approaches are related. PPO encourages a new stochastic policy to be close to a previous one in Kullback-Leibler distance. Whereas the term used in PPO modifies the objective function (and changes the set of stationary points), the proximal point algorithm optimizes and finds a stationary point of the original objective \mathcal{L} .

5 On the Difficulty of Offline Policy Evaluation and A new Benchmark

Evaluating precisely test policies on real systems is challenging as it would require making interventions on the system that are not possible. It is thus important to have reliable offline evaluation methods from logged propensities.

We achieve this by introducing a new evaluation protocol for stochastic policies with continuous actions, which ensures low false discoveries by using diagnostics and statistical testing. More precisely, we introduce (i) a large-scale, Criteo Off Policy Continuous Action (COPCAT) dataset collected through a massive controlled experiment in digital advertising and (ii) an evaluation protocol developed to assess the test performance of any policy while controlling false discoveries. While some previous work has considered simulated scenarios for evaluating continuous action policies [5, 16], to our knowledge, this is the first large-scale dataset with logged propensities where the data comes from a real-world decision-making system requiring continuous actions. From now on, we consider rewards instead of losses (negative losses), which are more meaningful in the applications considered.

The **Criteo Off Policy Continuous Action (COPCAT) dataset** comes from Criteo online advertising platform that ran an experiment involving a randomized, continuous policy for real-time bidding. Data has been properly anonymized so as to not disclose any private information. Each sample represents a bidding opportunity for which a context x in \mathbb{R}^d is observed, an action a in \mathbb{R}^+ is chosen according to a stochastic policy $\pi_0(a|x)$ that is logged along with the reward y in \mathbb{R} , the latter assumed to be caused partially by (x, a) . Particular care has been taken to guarantee that each sample $(x_i, a_i, \pi_0(a_i|x_i), y_i)$ is independent. As can be observed in summary statistics presented in Table 1(left), a typical feature of this application is the high variance of the reward, motivating the scale of the dataset to obtain precise counterfactual estimates. The link to download the dataset is available in the code repository: <https://github.com/criteo-research/optimization-continuous-action-crm>.

Table 1: (left) **Criteo Off Policy Continuous Action (COPCAT) dataset** summary statistics. (right) **Protocol** evaluation on synthetic data: correlation between online and off-policy estimates, false positive/negative rates for identifying improvements to the logging policy π_0 .

N	d	$\mathbb{E}[R]$	$\mathbb{V}[R]$	$\mathbb{V}[A]$	$\mathbb{P}(R \neq 0)$	corr.	FP rate	FN rate
120.10 ⁶	3	11.37	9455	.01	.07	.968	<1e-4	5e-4

The **evaluation protocol** is designed to obtain accurate counterfactual estimates for many possible policies and to decide if a new policy π is performing better than the logging π_0 with as few false positives as possible.

We found the approach described in Algorithm 1 to be successful in practice.

First, we use the self-normalized IPS (SNIPS) estimator \hat{R}_π to obtain lower variance estimates than IPS. Even though IPS may seem a natural choice, we evaluated this protocol on a pool of policies trained on the synthetic setup described in Sec. 6.1, which permits a comparison between online and off-policy estimates of test performance. We found that offline SNIPS estimates are highly correlated to the true (online) performance for a wide range of policies ($\rho = .968$, 30% higher than IPS).

Second, we detect unreliable estimates by leveraging importance sampling diagnostics based on the *effective sample size* n_{eff} , see [22], as well as the lower bound of a two-sided, $1 - \delta$ confidence interval for the

difference in performance ϵ_π (compared to π_0), which may be adjusted in case of multiple testing. We used the synthetic setup mentioned above to evaluate the importance of this filtering procedure, by measuring the ability of our method to detect improvements over the logging policy $\pi \succ \pi_0$ with known bad or good policies (i.e. sampled around π_0 and the analytical π^* , respectively). Interestingly, we found that a standard statistical significance test such as $\epsilon_\pi > 0$ was by itself no enough to guarantee low false positive rates. Nonetheless, adding a deviance criterion ν on the effective sample size significantly reduces the number of false positives, as reported in Table 1 (right). Such a diagnostic indeed helps to detect certain forms of overfitting to propensity scores in the SNIPS estimator [30] which can lead to invalid estimates. See Appendix 10 for more details.

Algorithm 1 Benchmark Evaluation Protocol

Input: $1 - \delta$: confidence of two-sided interval (def: .95); ν : a max deviance ratio (def: .01)

Output: mean counter-factual reward $\mathbb{E}_\pi[R]$ and decision $\pi \succ \pi_0$

1. Split dataset $\mathcal{D} \mapsto \mathcal{D}^{train}, \mathcal{D}^{valid}, \mathcal{D}^{test}$

2. Train new policy π on $\mathcal{D}^{train}(X, A, \pi_0(A|X), Y)$, tune π parameters on \mathcal{D}^{valid}

3. Estimate effective sample size diagnostic $n_{\text{eff}}(\pi, \mathcal{D}^{valid})$

4. Estimate $\hat{R}_\pi = \mathbb{E}_{\pi, test}^{\text{SNIPS}}[R]$ and lower bound $\epsilon_\pi = \lfloor \mathbb{E}_{\pi, test}^{\text{SNIPS}}[R] - \mathbb{E}_{\pi_0, test}[R] \rfloor_{1-\delta/2}$

Return \hat{R}_π if $\frac{n_{\text{eff}}}{n} > \nu$ else "invalid"; $\pi \succ \pi_0$ if $\frac{n_{\text{eff}}}{n} > \nu$ and $\epsilon_\pi > 0$

Note that Lefortier et al. [18] propose sanity checks for such offline evaluations that rely on estimating the expectation of importance sampling weights, notably by reporting empirical sums of importance weights when finding a counterfactual solution. However, they do not provide guidance on which threshold to use to reject policies. As such the question of how to compare policies is still unsolved. Instead, our protocol uses importance sampling diagnostics that discard solutions which do not meet a certain threshold that was tuned on synthetic cases for minimizing false discovery rates (see Table 1), both for validation when selecting hyperparameters and for test estimates.

6 Experiment Setup and Empirical Evaluation

We now introduce synthetic datasets that allow precise evaluation for any test policy; then, we present empirical findings on both these synthetic datasets and the Criteo Off Policy Continuous Action (COPCAT) dataset.

6.1 Synthetic and Semi-synthetic Settings

Potential prediction. We introduce simple synthetic environments with the following generative process: an unobserved random group index g in \mathcal{G} is drawn, which influences the drawing of a context x and of an unobserved “potential” p in \mathbb{R} , according to a joint conditional distribution $P_{X,P|G}$. The observed reward r is then a function of the context x , action a , and potential p . The causal graph corresponding to this process is given in Figure 1. Then, we generate three synthetic datasets (“noisymoos, noisycircles, and anisotropic”, illustrated in Appendix 11.1, Figure 13), with two-dimensional contexts on 2 or 3 groups and different choices of $P_{X,P|G}$.

The goal is then to find a policy $\pi(a|x)$ that maximizes the reward by adapting to the unobserved potential. For our experiments, potentials are normally distributed conditionally on the group index, $p|g \sim \mathcal{N}(\mu_g, \sigma^2)$. As many real-world applications feature a reward function that increases first with the action up to a peak and finally drops, we have chosen a piecewise linear function peaked at $a = p$ (see Appendix 11.1, Figure 14), that mimics overall reward buckets over the Criteo Off Policy Continuous Action (COPCAT) dataset presented in Section 5. In bidding applications, the potential may represent an unknown true value for an advertisement, and the reward is then maximized when the bid (the action) matches this value. In medicine, increasing drug




Figure 1: **Causal Graph of the synthetic setting.** A denotes the intervention action, X the feature context, G the unobserved group label, R the reward and P the unobserved potential.

dosage may increase treatment effectiveness but threshold effects soon appear and if dosage is pushed further, secondary effects may eclipse benefits [4].

Semi-synthetic setting with medical data. We follow the setup of [16] using a dataset on dosage of the Warfarin blood thinner drug [1]. The dataset consists of covariates about patients along with a dosage treatment prescription by a medical expert, which is a scalar value and thus makes the setting useful for continuous action modelling. While the dataset is supervised, we simulate a contextual bandit environment by using a hand-crafted reward function that is maximal for actions a that are within 10% of the expert’s therapeutic drug dosage, following [16]. See Appendix 11.1 for details.

6.2 Empirical Evaluation

We now present the empirical evaluation of our proposed CLP policy parameterization, as well as our optimization-related techniques. More extensive experiments and results are given in Appendix 12.

Evaluation methodology. In order to make our evaluation more realistic, we always perform hyperparameter selection by using off-policy estimates on a validation set with propensities obtained from the same logging policy as the training set, even for synthetic datasets. We also use the strategies outlined in Section 5 to discard policies which cannot be evaluated correctly on the validation set. The same strategy is used for test estimates on the Criteo Off Policy Continuous Action (COPCAT) dataset, while more accurate online estimates are used for test rewards in synthetic scenarios. See Appendix 11.2 for more details.

Counterfactual loss predictor (CLP). In Figure 2, we compare our CLP approach to a bucketing approach, which discretizes the action space and considers discrete-action policies with softmax parameterization [29] on the resulting buckets. By appropriately incorporating the continuous action structure, we can see that the CLP approach has much better performance than the discrete approach on the synthetic datasets, across all numbers of anchor points/buckets. The improvement is larger when using a small number of anchor points, illustrating that even a small number is sufficient to achieve precise action selection with our proposed Nyström approximation, in contrast to the discrete strategy which requires a much finer grid. We note that such a discretized approach requires ad-hoc propensity scores in order to work well, corresponding to the total mass of the bucket under the logging policy, which thus depends on the choice of buckets, and may impose additional burden and difficulties on the design of the logging system. The plots also show that the (stochastic) direct method has similar benefits in terms of robustness to anchor points, thanks to a similar Nyström parameterization. Nevertheless, it is overall outperformed by CLP, highlighting the benefits of using counterfactual methods compared to a direct fit to observed rewards, which may suffer from large bias.

Table 2 shows a comparison of test rewards with the scIPS estimator (other estimators are shown in Appendix 12.1). On the high-dimensional Warfarin dataset, following [16], we only consider the linear context parameterization baseline, for which [16] obtain a similar test reward, although our estimators and model selection criteria slightly differ. Overall, we find our new counterfactual loss predictor (CLP) parameterization to improve on all other choices of parameterizations on these datasets, which again highlights the effectiveness of the reward predictor at exploiting the continuous action structure.




Figure 2: **Continuous vs discrete.** Test rewards for CLP and (stochastic) direct method (SDM) with Nyström parameterization, versus a discrete approach, with varying numbers of anchor points. We add a minimal amount of noise to the deterministic DM in order to pass the n_{eff} validation criterion.

Table 2: **Comparison of policy parameterizations on all datasets for the (sc)IPS estimator.** Test rewards are shown with one standard deviation estimated across contexts.

	Noisycircles	NoisyMoons	Anisotropic	Warfarin	COPCAT
constant	0.6115 ± 0.0000	0.6116 ± 0.0000	0.6026 ± 0.0000	-10.050 ± 0.003	11.36 ± 0.13
linear	0.6113 ± 0.0001	0.7326 ± 0.0001	0.7638 ± 0.0005	-10.291 ± 0.004	8.00 ± 0.72
non-linear	0.6959 ± 0.0001	0.7281 ± 0.0001	0.7448 ± 0.0008	-	8.78 ± 0.51
clp	0.7674 ± 0.0008	0.7805 ± 0.0004	0.7703 ± 0.0002	-9.988 ± 0.001	11.44 ± 0.10

Optimization-driven approaches. Figure 3 shows the improvements in test reward and in training objective of our optimization-driven strategies, namely the soft-clipping estimator and the use of the proximal point algorithm. Figure 3 (left) illustrates the benefits of the proximal point method when optimizing the (non-convex) CRM objective in a wide range of hyperparameter configurations, and Figure 3 (center) shows that in many cases this results in an improvement in test reward as well. Here, each point compares the test metric for fixed models as well as initialization seeds, while optimizing the remaining hyperparameters on the validation set. In Figure 3 (right), the points correspond to different choices of the clipping parameter M , models and initialization, with the rest of the hyper-parameters optimized on the validation set. This plot also shows that the smoothness of the objective in soft clipping compared to hard clipping provides benefits in terms of test reward, perhaps thanks to a more favorable optimization landscape. Overall, these figures confirm that the optimization perspective is an important one when considering CRM problems.




Figure 3: **Optimization-driven approaches.** Relative improvements in the training objective from using the proximal point method (left), comparison of test rewards for proximal point vs the simpler gradient-based method (center), and for soft- vs hard-clipping (right).

7 Conclusion

In this work, we focus on learning stochastic continuous policies with parametric contextual modellings. We first propose a new contextual estimator with rich expressive power through the joint modelling of contexts and actions, showing superior empirical performance on a wide range of problems. Second, we underline the importance of optimization in CRM formulations with a soft-clipping estimator and the use of proximal point methods. Third, whereas the community has focused on synthetic or semi-synthetic setups, we propose an offline evaluation protocol and a new large-scale dataset, which, to the best of our knowledge, is the first to be based on real-world logged propensities.

Acknowledgments

The authors thank Christophe Renaudin (Criteo AI Lab) for his help in the data collection and engineering which was necessary for this project. JM was supported by the ERC grant number 714381 (SOLARIS project) and by ANR 3IA MIAI@Grenoble Alpes, (ANR-19-P3IA-0003). AB acknowledges support from the ERC grant number 724063 (SEQUOIA project).

References

- [1] Estimation of the Warfarin dose with clinical and pharmacogenetic data. *New England Journal of Medicine*, 360(8):753–764, 2009.
- [2] A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, and R. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning (ICML)*, 2014.
- [3] Z. Ahmed, N. Le Roux, M. Norouzi, and D. Schuurmans. Understanding the impact of entropy on policy optimization. In *International Conference on Machine Learning (ICML)*, 2019.
- [4] C. D. Barnes and L. G. Eltherington. *Drug dosage in laboratory animals: a handbook*. Univ of California Press, 1966.
- [5] D. Bertsimas and C. McCord. Optimization over continuous and multi-dimensional decisions with observational data. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [6] L. Bottou, J. Peters, J. Quiñero Candela, D. X. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, and E. Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research (JMLR)*, 14(1):3207–3260, Jan. 2013.
- [7] M. Demirer, V. Syrgkanis, G. Lewis, and V. Chernozhukov. Semi-parametric efficient policy learning with continuous actions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [8] M. Dudik, J. Langford, and L. Li. Doubly robust policy evaluation and learning. In *International Conference on Machine Learning (ICML)*, 2011.
- [9] D. J. Foster and V. Syrgkanis. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019.
- [10] M. Fukushima and H. Mine. A generalized proximal point algorithm for certain non-convex minimization problems. *International Journal of Systems Science*, 12(8):989–1000, 1981.
- [11] K. Hirano and G. W. Imbens. The propensity score with continuous treatments. *Applied Bayesian modeling and causal inference from incomplete-data perspectives*, 226164:73–84, 2004.
- [12] D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.

- [13] N. Jiang and L. Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2016.
- [14] T. Joachims, A. Swaminathan, and M. de Rijke. Deep learning with logged bandit feedback. In *International Conference on Learning Representations (ICLR)*, 2018.
- [15] S. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2002.
- [16] N. Kallus and A. Zhou. Policy evaluation and optimization with continuous treatments. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- [17] J. Langford and T. Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in Neural Information Processing Systems (NIPS)*, 2008.
- [18] D. Lefortier, A. Swaminathan, X. Gu, T. Joachims, and M. de Rijke. Large-scale validation of counterfactual learning methods: A test-bed. 2016. URL <http://arxiv.org/abs/1612.00367>.
- [19] L. Li, W. Chu, J. Langford, T. Moon, and X. Wang. An unbiased offline evaluation of contextual bandit algorithms with generalized linear models. In *Proceedings of the Workshop on On-line Trading of Exploration and Exploitation 2*, 2012.
- [20] D. C. Liu and J. Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- [21] A. Maurer and M. Pontil. Empirical bernstein bounds and sample variance penalization. In *Conference on Learning Theory (COLT)*, 2009.
- [22] A. B. Owen. *Monte Carlo theory, methods and examples*. 2013.
- [23] C. Paquette, H. Lin, D. Drusvyatskiy, J. Mairal, and Z. Harchaoui. Catalyst for gradient-based nonconvex optimization. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- [24] J. M. Robins and A. Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.
- [25] R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- [26] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [27] Y. Su, L. Wang, M. Santacatterina, and T. Joachims. Cab: Continuous adaptive blending for policy evaluation and learning. In *International Conference on Machine Learning*, pages 6005–6014, 2019.
- [28] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems (NIPS)*, 2000.
- [29] A. Swaminathan and T. Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning (ICML)*, 2015.
- [30] A. Swaminathan and T. Joachims. The self-normalized estimator for counterfactual learning. In *Advances in Neural Information Processing Systems (NIPS)*. 2015.
- [31] Y.-X. Wang, A. Agarwal, and M. Dudík. Optimal and adaptive off-policy evaluation in contextual bandits. In *International Conference on Machine Learning (ICML)*, 2017.

- [32] C. K. Williams and M. Seeger. Using the nystrom method to speed up kernel machines. In *Adv. Neural Information Processing Systems (NIPS)*, 2001.
- [33] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

Appendix

This appendix is organized as follows: in Appendix 8, we present a review of off-policy estimators, then in Appendix 9, we present discussions and toy experiments to motivate the need for clipping strategies on real datasets. Next, we provide motivations for the offline evaluation protocol with experiments justifying the need for appropriate diagnostics and statistical testing for importance sampling. Then, Appendix 11 is devoted to experimental details that were omitted from the main paper for space limitation reasons, and which are important for reproducing our results (see also the code provided with the submission). In Appendix 12, we present additional experimental results to those in the main paper. Finally, we provide proof for the Proposition 4.1 in Appendix 13.

8 Review of Off-policy Estimators

In Equation (3), the counterfactual approach tackles the distribution mismatch between the logging policy $\pi_0(\cdot|x)$ and a policy π in Π via importance sampling and performs inverse propensity scoring (IPS, [12]). However, the empirical (IPS) estimator has large variance and may overfit negative feedback values y_i for samples that are unlikely under π_0 (see motivation for clipped estimators in 9.2), resulting in higher variances. Clipping the importance sampling weights in Eq. (8) as in [6] mitigates this problem, leading to a new clipped (cIPS) estimator:

$$\hat{R}_{\text{cIPS}}^M(\pi) = \frac{1}{n} \sum_{i=1}^n y_i \min \left\{ \frac{\pi(a_i|x_i)}{\pi_{0,i}}, M \right\}. \quad (8)$$

Smaller values of M reduce the variance of $\hat{R}^M(\pi)$ but induce a larger bias. Swaminathan and Joachims [29] also propose adding an empirical variance penalty term controlled by a factor λ to the empirical risk $\hat{R}^M(\pi)$, leading to a regularized objective with hyperparameters M and λ for clipping and variance penalization, respectively.

Swaminathan and Joachims [30] also introduce a regularization mechanism for tackling the so-called *propensity overfitting* issue, occurring with rich policy classes, where the method would focus only on maximizing (resp. minimizing) the sum of ratios $\pi(a_i|x_i)/\pi_{0,i}$ for negative (resp. positive) losses. This effect is corrected through the following *self-normalized importance sampling* (SNIPS) estimator (see also [22]), which is equivariant to additive shifts in loss values:

$$\hat{R}_{\text{SNIPS}}(\pi) = \frac{\sum_{i=1}^n y_i w_i^\pi}{\sum_{i=1}^n w_i^\pi}, \quad \text{with } w_i^\pi = \frac{\pi(a_i|x_i)}{\pi_{0,i}}. \quad (9)$$

Direct methods (DM) fit the loss values over contexts and actions in observed data with an estimator $\hat{\eta}(x, a)$, for instance by using ridge regression to fit $y_i \approx \hat{\eta}(x_i, a_i)$, and to then use the deterministic greedy policy $\hat{\pi}_{\text{DM}}(x) = \arg \min_a \hat{\eta}(x, a)$. These may suffer from large bias since it focuses on estimating losses mainly near actions that appear in the logged data but have the benefit of avoiding the high-variance problems of IPS-based methods. While such greedy deterministic policies may be sufficient for exploitation, stochastic policies may be needed in some situations, for instance when one wants to still encourage some exploration in a future round of data logs, perhaps for non-stationarity issues. Using a stochastic policy also allows us to obtain more accurate off-policy estimates when performing cross-validation on logged data. Then, we consider a stochastic version of the direct method by adding some Gaussian noise with variance σ^2 :

$$\hat{\pi}_{\text{SDM}}(\cdot|x) = \mathcal{N}(\hat{\pi}_{\text{DM}}(x), \sigma^2), \quad (10)$$

In the context of offline evaluation on bandit data, such a smoothing procedure may also be seen as a form of kernel smoothing for better estimation [16].

Additionally, such direct loss estimators can be effective when few samples are available, and may be combined with IPS estimators in the so-called doubly robust estimator (DR, see, e.g. [8]). This approach consists of correcting the bias of the DM estimator by applying IPS to the residuals $y_i - \hat{\eta}(x_i, a_i)$, thus using $\hat{\eta}$ as a control variate to decrease the variance of IPS.

9 Motivation for Clipped Estimators

In this section we provide an illustration of the logarithmic soft clipping and a motivation example for clipping strategies in counterfactual systems.

9.1 Soft clipping details

Soft clipping strategies are useful for learning on real-life datasets with outliers and to avoid variance in IPS estimates. The lower the clipping parameter is, the more biased the estimate is, but the higher it is the more likely it is to overfit, particularly for low propensities. We illustrate in Fig. 4 the expression of the logarithmic clipping:

$$\zeta(w_i, M) = \begin{cases} w_i & \text{if } w_i \leq M \\ \alpha(M) \log(w_i + \alpha(M) - M) & \text{otherwise,} \end{cases} \quad (11)$$

where $\alpha(M)$ is such that $\alpha(M) \log(\alpha(M)) = M$.




Figure 4: Soft Clipping of importance sampling weight w .

9.2 A Toy Example for Soft Clipping

In Figure 5 we provide an example of unbounded variance and loss overfitting problem.

We recall the data generation: a hidden group label g in \mathcal{G} is drawn, and influences the associated context distribution x and of an unobserved potential p in \mathbb{R} , according to a joint conditional distribution $P_{X,P|G}$. The observed reward r is then a function of the context x , action a , and potential p . Here, we design one outlier (big red dark dot on Figure 4 left). This point has a noisy reward r , higher than neighbors, and a potential p high as its neighbors have a low potential. We artificially added a noise in the reward function f that can be written as:

$$r = f(a, x, p) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1)$$

As explained in Section 6.1, the reward function is a linear function, with its maximum localized at the point $x = p(x)$, i.e. at the potential sampled. The observability of the potential p is only through this reward function f . Hereafter, we compare the optimal policy computed, using different type of estimators.

The task is to predict the high potentials (red circles) and low potentials (blue circles) in the ground truth data (left). Unfortunately, a rare event sample with high potential is put in the low potential cluster (big dark red dot). The action taken by the logging policy is low while the reward is high: this sample is an

outlier because it has a high reward while being a high potential that has been predicted with a low action. The resulting unclipped estimator is biased and overfits this high reward/low propensity sample. The rewards of the points around this outlier are low as the diameter of the points in the middle figure show. Inversely, clipped estimator with soft-clipping succeeds to learn the potential distributions, does not overfit the outlier, and has larger rewards than the clipping policies as the diameter of the points show.




Figure 5: **Unbounded variance and loss overfitting.** Unlikely ($\pi_{0,i} \approx 0$) sample $(x_1, x_2) = (0.6, 0.)$ with high reward r (left) results in larger variance and loss overfitting for the unclipped estimator (middle) unlike clipped estimator (right).

10 Motivation for Offline Evaluation Protocol

In this part we demonstrate the offline/online correlation of the estimator we use for real-world systems and for validation of our methods even in synthetic and semi-synthetic setups. We provide further explanations of the necessity of importance sampling diagnostics and we perform experiments to empirically assess the rate of false discoveries of our protocol.

10.1 Correlation of Self-Normalized Importance Sampling with Online Rewards

We show in Figures 6,8,7 comparisons of IPS and SNIPS against an on-policy estimate of the reward for policies obtained from our experiments for linear and non-linear contextual modellings on the synthetic datasets, where policies can be directly evaluated online. Each point represents an experiment for a model and a hyperparameter combination. We measure the R^2 score to assess the quality of the estimation, and find that the SNIPS estimator is indeed more robust and gives a better fit to the on-policy estimate. Note also that overall the IPS estimates illustrate severe variance compared to their SNIPS estimate. While SNIPS indeed reduces the variance of the estimate, the bias it introduces does not deteriorate too much its (positive) correlation with the online evaluation.

These figures further justify the choice of the self-normalized estimator SNIPS [30] for offline evaluation and validation to estimate the reward on held-out logged bandit data. The SNIPS estimator is indeed more robust to the reward distribution thanks to its equivariance property to additive shifts and does not require hyperparameter tuning.

10.2 Importance Sampling Diagnostics

Importance sampling estimators rely on weighted observations to address the distribution mismatch for offline evaluation, which may cause large variance of the estimator. Notably, when the evaluated policy differs too much from the logging policy, many importance weights are large and the estimator is inaccurate. We provide here a motivating example to illustrate the effect of importance sampling diagnostics in a simple scenario.

When evaluating with SNIPS, we may consider an “effective sample size” quantity given in terms of the importance weights $w_i = \pi(a_i|x_i)/\pi_0(a_i|x_i)$ by

$$n_e = (\sum_{i=1}^n w_i)^2 / \sum_{i=1}^n w_i^2.$$




Figure 6: **Correlation between offline and online estimates on Anisotropic synthetic data.** Linear (left) and non-linear (right) contextual modellings. Ideal fit would be $y = x$.




Figure 7: **Correlation between offline and online estimates on NoisyCircles synthetic data.** Linear (left) and non-linear (right) contextual modellings. Ideal fit would be $y = x$.

When this quantity is much smaller than the sample size n , this indicates that only few of the examples contribute to the estimate, so that the obtained value is likely a poor estimate. Apart from that, we note also that IPS weights have an expectation of 1 when summed over the logging policy distribution (that is $\mathbb{E}_{(x,a) \sim \pi_0} [\pi(a|x)/\pi_0(a|x)] = 1$). Therefore, another sanity check, which is valid for any estimator, is to look for the empirical mean $1/n \sum_{i=1}^n \pi(a_i|x_i)/\pi_{0,i}$ and compare its deviation to 1. In the example below, we illustrate three diagnostics: (i) the one based on effective sample size described in Section 5; (ii) confidence intervals, and (iii) empirical mean of IPS weights. The three of them coincide and allow us to remove test estimates when the diagnostics fail.

Example 10.1. *What-if simulation:* For x in \mathbb{R}^d , let $\max(x) = \max_{1 \leq j \leq d} x_j$; we wish to estimate $\mathbb{E}(\max(X))$ for X i.i.d $\sim \pi_\mu = \mathcal{N}(\mu, \sigma)$ where samples are drawn from a logging policy $\pi_0 = \log \mathcal{N}(\lambda_0, \sigma_0)$ ($d = 3, (\lambda_0, \sigma_0) = (1, 1/2)$) and analyze parameters μ around the mode of the logging policy μ_{π_0} with fixed variance $\sigma = 1/2$. In this parametrized policy example, we see in Fig. 9 that $n_e/n \ll 1$, confidence interval range increases and $\sum_{i=1}^n \frac{\pi(a_i|x_i)}{\pi_{0,i}} \neq 1$ when the parameter μ of the policy being evaluated is far away from the logging policy mode μ_{λ_0} .

Note that in this example, the parameterized distribution that is learned (multivariate Gaussian) is not the same as the parameterized distribution of the logging policy (multivariate Lognormal) which skewness may explain the asymmetry of the plots. This points out another practical problem: even though different parametrization of policies is theoretically possible, the probability density masses overlap is in practice what is most important to ensure successful importance sampling. This observation is of utmost interest for real-life applications where the initialization of a policy to be learned needs to be "close" to the logging policy; otherwise importance sampling may fail from the very first iteration of an optimization in learning problems.

10.3 Experimental validation of the protocol

While the goal of counterfactual learning is to find a policy $\hat{\pi}$ which is as close a possible to the optimal policy π^* , importance sampling estimates are somehow limited in practice when the optimal policy π^* is too




Figure 8: **Correlation between offline and online estimates on NoisyMoons synthetic data.** Linear (left) and non-linear (right) contextual modellings. Ideal fit would be $y = x$.




Figure 9: **Importance sampling diagnostics.** Ideal importance sampling: i) effective sample n_e/n close to 1, ii) low confidence intervals (C.I.s) for \hat{m} , iii) empirical mean $\frac{1}{n} \sum_i w_i$ close to 1. Note that when μ differs too much from μ_{π_0} , importance sampling fails.

"far" from the logging policy π_0 . Therefore for real-world applications one may focus on improving over the logging policy to deploy a policy in production systems in order to progressively collect new data. Thus, it is required to derive level of confidences when estimating the risk of an empirically optimized policy in order to make the decision " $\hat{\pi} \succeq \pi_0$ ".

We empirically evaluate our offline evaluation protocol on a toy setting where we compare the predictions made with the IPS offline metric as well as the SNIPS offline metric. Offline evaluations of policies $\hat{\pi}$ illustrated in Fig. (10) are estimated from logged data $(x_i, a_i, y_i, \pi_0)_{i=1\dots n}$ where $a_i \sim \pi_0(\cdot|x_i)$ and where rewards r would be optimal for the oracle policy π^* .

Specifically, we evaluate the quality/efficiency of offline estimates for policies (i) "close" to the logging policy and (ii) "close" to the oracle optimal policy. For both setups (i) and (ii), we compare the number of false positives and false negatives of the two offline protocols for 2000 initializations and also show histograms of the differences between online and offline boundary decisions for " $\hat{\pi} \succeq \pi_0$ " which is the lower bound $\epsilon_\pi = \lfloor \mathbb{E}_{\pi, test}^{SNIPS}[R] - \mathbb{E}_{\pi_0, test}[R] \rfloor_{1-\delta/2}$ as referred in Algorithm 1.

10.3.1 Perturbation to the logging policy π_0

Knowing the closed form of the logging policy π_0 and its parameters, we add Gaussian noise to the parameters of π_0 and use the offline evaluation protocol to discard solutions with importance sampling diagnostics (effective sample sizes < 0.01) and estimate confidence intervals with bootstrap procedures on the test set.

We observe that the SNIPS estimate has fewer false positives and false negatives, as shown in Table 3. Histograms of boundary decisions $\epsilon_\pi = \lfloor \mathbb{E}_{\pi, test}^{SNIPS}[R] - \mathbb{E}_{\pi_0, test}[R] \rfloor_{1-\delta/2}$ illustrated in Fig. 11 show that the IPS estimate underestimates the value of the reward with regard to the online estimate and has a lot of variance, unlike the SNIPS estimates which exhibits less variance and seems centered around the online estimate.




Figure 10: **Offline evaluation protocol.**

Table 3: **Comparison of false positives and false negatives:** Perturbation to the logging policy π_0

	Offline Protocol	IPS		SNIPS	
		$\hat{\pi} \succeq \pi_0$	Fail	$\hat{\pi} \succeq \pi_0$	Fail
"Truth"	$\hat{\pi} \succeq \pi_0$	1282	24	1296	10
	Fail	19	675	0	694

10.3.2 Perturbation to the oracle policy π^*

Knowing the closed form of the oracle policy π^* and its parameters, we add Gaussian noise to its parameters and use the same protocol as before. This setup is "harder" than the previous one since importance sampling is more likely to fail when the evaluated policy is too "far" from the logging policy.

We observe that the SNIPS estimate has a drastically lower number of false negatives than the IPS estimate, though it has more false positives as shown in Table 4. Once again as illustrated in Figure 12, the IPS estimate underestimates the value of the reward with regard to the online estimate and large variance which may explain why IPS does not have false positives where the evaluated policy is far from the logging policy.

Table 4: **Comparison of false positives and false negatives:** perturbations to the optimal policy π^*

	Offline Protocol	IPS		SNIPS	
		$\hat{\pi} \succeq \pi_0$	Fail	$\hat{\pi} \succeq \pi_0$	Fail
"Truth"	$\hat{\pi} \succeq \pi_0$	1565	67	1631	1
	Fail	0	368	6	362

11 Details on the Experiment Setup and Reproducibility

In this section we give additional details on synthetic and semi-synthetic datasets, we provide details on the evaluation methodology and information for experiment reproducibility.




Figure 11: **Histogram of differences between online reward and offline lower confidence bound.** Perturbation to the logging policy π_0




Figure 12: **Histogram of differences between online reward and offline lower confidence bound.** Perturbation to the optimal policy π^*

11.1 Synthetic and Semi-Synthetic setups

Synthetic setups The three generated synthetic datasets called “noisymoons, noisycircles, and anisotropic”, are illustrated in Figure 13, with two-dimensional contexts $\mathcal{X} = \mathbb{R}^2$ and 2 or 3 groups.

As many real-world applications feature a reward function that increases first with the action, then plateaus and finally drops, we have chosen a piecewise linear function as shown in Fig. 14 that mimics reward buckets over the Criteo Off Policy Continuous Action (COPCAT) dataset presented in Section 5.

Semi-synthetic medical setup The semi-synthetic loss inputs prescriptions from medical experts to obtain $y(a, x) = \max(|a - t^*| - 0.1t^*, 0)$. The logging policy π_0 samples actions $a \sim \pi_0$ contextually to a patient’s BMI z -score $Z_{BMI} = \frac{x_{BMI} - \mu_{BMI}}{\sigma_{BMI}}$ and can be analytically written with i.i.d noise $e \sim \mathcal{N}(0, 1)$, moments of the therapeutic dose distribution μ_T^*, σ_T^* such that $a = \mu_T^* + \sigma_T^* \sqrt{\theta} Z_{BMI} + \sigma_T^* \sqrt{1 - \theta} e$ ($\theta = 0.5$ in the setup of [16]). The logging probability density function thus is a continuous density of a standard normal distribution over the quantity $\frac{a - \mu_T^* + \sigma_T^* \sqrt{\theta} Z_{BMI}}{\sigma_T^* \sqrt{1 - \theta}}$.




Figure 13: **Contexts (points in \mathbb{R}^2), and potentials represented by a color map for the synthetic datasets.** Learned policies should vary with the context to adapt to the underlying potentials.




Figure 14: **Synthetic reward engineering.** The synthetic reward (left) is inspired from real-dataset reward buckets (right).

11.2 Evaluation methodology.

For synthetic datasets, we generate training, validation, and test sets of size 10000 each. For the Criteo Off Policy Continuous Action (COPCAT) dataset, we consider a 50%-25%-25% training-validation-test sets. We then run each method with 5 different random initializations such that the initial policy is close to the logging policy. Hyperparameters are selected on a validation set with logged bandit feedback, by using an offline SNIPS estimate of the obtained policies, while discarding solutions deemed unsafe with the importance sampling diagnostic. For estimating the final test performance and confidence intervals on synthetic datasets, we use an online estimate by leveraging the known reward function and taking a Monte Carlo average with 100 action samples per context: this accounts for the randomness of the policy itself over given fixed samples. For offline estimates we leverage the randomness across samples to build confidence intervals: we use a 100-fold bootstrap and take percentiles of the distribution of rewards as given in Algo 1. For the Criteo Off Policy Continuous Action (COPCAT) dataset, we use the latter methodology in our reported test metrics.

11.3 Reproducibility

We provide code for reproducibility and all experiments were run on a cpu cluster, each node consisting on 24 CPU cores (2 x Intel(R) Xeon(R) Gold 6146 CPU@ 3.20GHz), with 500Go of RAM.

Policy parametrization. In our experiments, we consider two forms of parametrizations: (i) a lognormal distribution with $\theta = (\theta_\mu, \sigma)$, $\pi_{(\mu, \sigma)} = \log \mathcal{N}(m, s)$ with $s = \sqrt{\log(\sigma^2/\mu^2 + 1)}$; $m = \log(\mu) - s^2/2$, so that $\mathbb{E}_{a \sim \pi_{(\mu, \sigma)}}[a] = \mu$ and $\text{Var}_{a \sim \pi_{(\mu, \sigma)}}[a] = \sigma^2$; (ii) a normal distribution $\pi_{(\mu, \sigma)} = \mathcal{N}(\mu, \sigma)$. In both cases, the mean μ may depend on the context (see Section 4), while the standard deviation σ is a learned constant. We add a positivity constraint for σ and add an entropy regularization term to the objective in order to encourage exploratory policies and avoid degenerate solutions.

Models. For our experiments involving anchor points, we took $m = 10$ anchor points on the synthetic and semi-synthetic dataset and $m = 3$ for the Criteo Off Policy Continuous Action (COPCAT) dataset. For

parametrized distributions, we experimented both with normal and lognormal distributions on all datasets, and different baseline parameterizations including constant, linear and quadratic. Some of our experiments on low-dimensional datasets also consider a stratified piece-wise contextual parameterization, which partitions the space by bucketizing each feature by taking K quantiles, with $K = 4$, and taking the cross product of these partitions for each feature.

Hyperparameters. In Table 5 we show the hyperparameters considered to run the experiments to reproduce all the results. Note that the grid of hyperparameters is larger for synthetic data.

Table 5: Table of hyperparameters for the Synthetic, Warfarin and Criteo Off Policy Continuous Action (COPCAT) datasets

	Synthetic	Warfarin	COPCAT
Variance reg. λ	{0., 0.001, 0.01, 0.1, 1, 10, 100}	{0.00010.0010.010.1}	{0., 0.001, 0.1}
Clipping M	{1, 1.7, 2.8, 4.6, 7.7, 12.9, 21.5, 35.9, 59.9, 100.0}	{1, 2.1, 4.5, 9.5, 20}	{1, 2.1, 4.5, 9.5, 10, 20, 100}
Prox. κ	{0.001, 0.01, 0.1, 1}	{0.001, 0.01, 0.1}	{0.001, 0.01, 0.1}
Reg. param. C	{0.00001, 0.0001, 0.001, 0.01, 0.1}	{0.00001, 0.0001, 0.001, 0.01, 0.1}	{0.00001, 0.0001, 0.001, 0.01, 0.1}
Reg. entropy ϵ	{0.0001, 0.01, 1.}	{0.0001, 0.01}	{0.0001, 0.01}
Softmax γ	{1, 10, 100}	{1, 10}	{0.1, 0.5, 1}

12 Additional Results and Additional Evaluation Metrics

In this section we provided additional results on both contextual modeling and optimization driven approaches of CRM.

12.1 Contextual Modeling Comparisons

We show additional comparisons of contextual models for stochastic policies for other estimators in Table 6. On the synthetic datasets, it is easy to beat the logging policy with high level of confidence and with the importance sampling diagnostics. For the Criteo Off Policy Continuous Action (COPCAT) dataset, we build offline statistical confidence interval with bootstrap percentiles as described in Algo 1.

For the Criteo Off Policy Continuous Action (COPCAT) dataset, it is highly difficult to obtain policies with high offline estimates that verify the importance sampling diagnostics with bootstrap confidence intervals that do not intersect with the logging policy, but note that a significant improvement is achieved for the CLP model with the SNIPS estimator (see the asterisk in Table 6, which indicates a significant improvement over the logging policy according to our protocol in Algorithm 1). We also note that the SNIPS estimator in CRM achieves the best performances for both the semi-synthetic and COPCAT datasets, while for the synthetic datasets, IPS and clipped IPS seem to lead to better policies.

12.2 Optimization Driven Approaches of CRM

In this part we provide additional results on optimization driven approaches of CRM for the Noisycircles, Anisotropic, Warfarin and Criteo Off Policy Continuous Action (COPCAT) datasets.

Both Noisycircles and Anisotropic datasets in Figure 15 show the improvements in test reward and in training objective of our optimization-driven strategies, namely the soft-clipping estimator and the use of the proximal point algorithm. Overall we see that for most configurations, the proximal point method better optimizes the objective function and provides better test performances, while the soft-clipping estimator performs better than its hard-clipping variant, which may be attributed to the better optimization properties. For semi-synthetic Warfarin and real-world COPCAT datasets in Figure 15 we also show the improvements in test reward and in training objective of our optimization-driven strategies. More particularly we demonstrate the effectiveness of proximal point methods on the Warfarin dataset where most proximal configurations

Table 6: **Comparison of policy parameterizations on all datasets for different CRM estimators.** Test rewards are shown with one standard deviation estimated across contexts. This is a more comprehensive version of Table 2.

		Noisycircles	NoisyMoons	Anisotropic	Warfarin	COPCAT
Logging policy π_0		0.5301	0.5301	0.4533	-13.3769	11.34
IPS	constant	0.6114 \pm 0.0001	0.6114 \pm 0.0001	0.6018 \pm 0.0001	-10.066 \pm 0.003	11.35 \pm 0.01
	linear	0.6115 \pm 0.0001	0.7329 \pm 0.0001	0.7465 \pm 0.0007	-10.295 \pm 0.009	11.37 \pm 0.02
	non-linear	0.6943 \pm 0.0001	0.7137 \pm 0.0001	0.7301 \pm 0.0010	-	10.17 \pm 0.12
	clp	0.7663 \pm 0.0003	0.7432 \pm 0.0004	0.8047 \pm 0.0002	-11.392 \pm 0.003	11.35 \pm 0.03
cIPS	constant	0.6116 \pm 0.0001	0.6116 \pm 0.0001	0.6026 \pm 0.0001	-10.064 \pm 0.003	11.36 \pm 0.13
	linear	0.6115 \pm 0.0001	0.7337 \pm 0.0001	0.7465 \pm 0.0007	-9.2159 \pm 0.030	11.35 \pm 0.02
	non-linear	0.6927 \pm 0.0001	0.7287 \pm 0.0001	0.7483 \pm 0.0003	-	10.17 \pm 0.08
	clp	0.7925 \pm 0.0006	0.7940 \pm 0.0004	0.8011 \pm 0.0002	-10.007 \pm 0.003	11.34 \pm 0.09
scIPS	constant	0.6115 \pm 0.0000	0.6116 \pm 0.0000	0.6026 \pm 0.0000	-10.050 \pm 0.003	11.36 \pm 0.13
	linear	0.6113 \pm 0.0001	0.7326 \pm 0.0001	0.7638 \pm 0.0005	-10.291 \pm 0.004	8.00 \pm 0.72
	non-linear	0.6959 \pm 0.0001	0.7281 \pm 0.0001	0.7448 \pm 0.0008	-	8.78 \pm 0.51
	clp	0.7674 \pm 0.0008	0.7805 \pm 0.0004	0.7703 \pm 0.0002	-9.988 \pm 0.001	11.44* \pm 0.10
SNIPS	constant	0.6115 \pm 0.0001	0.6115 \pm 0.0001	0.5930 \pm 0.0001	-8.92 \pm 0.0006	11.32 \pm 0.13
	linear	0.6115 \pm 0.0001	0.7360 \pm 0.0001	0.7103 \pm 0.0003	-14.158 \pm 0.0003	10.34 \pm 0.12
	non-linear	0.6969 \pm 0.0001	0.7370 \pm 0.0001	0.5801 \pm 0.0002	-	11.13 \pm 0.08
	clp	0.6972 \pm 0.0001	0.74091 \pm 0.0004	0.7899 \pm 0.0002	-8.505 \pm 0.007	11.48* \pm 0.14

perform better than the base algorithm. Moreover, soft-clipping strategies perform better than its hard-clipping variant on real-world dataset with outliers and noises, which demonstrate the effectiveness of this smooth estimator for real-world setups.

Note also that using the importance sampling diagnostics is crucial to ensure an estimator is reliable on the test set. Indeed, when removing the importance sampling diagnostics, experiments show results which are either drastically below or much higher than the logging policy. The importance sampling diagnostics thus helps avert unsuccessful estimates and prevents from overestimating possible solutions. On the open dataset especially, beating the logging policy with a high level of confidence and makes offline estimation on real-life bandit feedback datasets a fundamental problem.

13 Generalization Bound for the Soft-Clipping Estimator \hat{R}_{scIPS}^M (proof of Proposition 4.1)

In this section, we derive a simple generalization bound on the risk for the soft-clipping estimator with empirical variance regularization, following Maurer and Pontil [21], Swaminathan and Joachims [29]. Recall the estimator

$$\hat{R}_{\text{scIPS}}^M(\pi) = \frac{1}{n} \sum_{i=1}^n \nu_i(\pi), \quad \text{with } \nu_i(\pi) = y_i \zeta \left(\frac{\pi(a_i|x_i)}{\pi_0(a_i|x_i)}, M \right),$$

and the empirical variance which is used for regularization:

$$\hat{V}(\pi) = \frac{1}{n-1} \sum_{i=1}^n (\nu_i(\pi) - \bar{\nu}(\pi))^2, \quad \text{with } \bar{\nu}(\pi) = \frac{1}{n} \sum_{i=1}^n \nu_i(\pi).$$

We assume that losses $y_i \in [-1, 0]$ almost surely, as in [29], and make the additional assumption that the importance weights $\pi(a_i|x_i)/\pi_0(a_i|x_i)$ are upper bounded by a constant W almost surely for all $\pi \in \Pi$. This is verified, for instance, if all policies have a given compact support (e.g., actions are constrained to belong to a given interval) and π_0 puts mass everywhere in this support.

We state the bound for a finite policy class for simplicity, but we remark below that it extends to infinite policy classes through covering numbers. Note that while the bound requires importance weights bounded by a constant W , the bound only scales logarithmically with W when $W \gg M$.

Proposition 13.1 (Generalization bound for finite policy class Π). *Assume losses in $[-1, 0]$ and importance weights bounded by W . With probability $1 - \delta$, we have, for all $\pi \in \Pi$,*

$$R(\pi) \leq \hat{R}_{scIPs}^M(\pi) + \sqrt{\frac{2\hat{V}(\pi) \log(2|\Pi|/\delta)}{n}} + S \frac{7 \log(2|\Pi|/\delta)}{3(n-1)}, \quad (12)$$

where $S = \zeta(W, M) = O(\log W)$.

Proof. Let

$$f_i(\pi) = 1 + \frac{y_i}{S} \zeta \left(\frac{\pi(a_i|x_i)}{\pi_0(a_i|x_i)}, M \right).$$

We have $f_i(\pi) \in [0, 1]$ and $\mathbb{E}_{(x_i, a_i, y_i) \sim \mathcal{D}_{\pi_0}}[f_i(\pi)] = 1 + R^M(\pi)/S$, with

$$R^M(\pi) = \mathbb{E}_{(x, a, y) \sim \mathcal{D}_{\pi_0}} \left[y \zeta \left(\frac{\pi(a|x)}{\pi_0(a|x)}, M \right) \right].$$

We can now apply the concentration bound of Maurer and Pontil [21, Corollary 5] to the f_i and rescale appropriately by S to obtain that with probability $1 - \delta$, for all $\pi \in \Pi$,

$$R^M(\pi) \leq \hat{R}_{scIPs}^M(\pi) + \sqrt{\frac{2\hat{V}(\pi) \log(2|\Pi|/\delta)}{n}} + S \frac{7 \log(2|\Pi|/\delta)}{3(n-1)}.$$

We then conclude by noting that $R(\pi) \leq R^M(\pi)$, since $\ell(a) \leq 0$ and $\zeta(w, M) \leq w$ for all w . □

We note that the result above can be extended to infinite policy classes by essentially replacing $|\Pi|$ above with an ℓ_∞ covering number of the set $\{(f_1(\pi), \dots, f_n(\pi)) : \pi \in \Pi\}$, by leveraging Maurer and Pontil [21, Theorem 6] as in [29].




Figure 15: **Optimization-driven approaches** (NoisyCircles, Anisotropic, Warfarin and Criteo Off Policy Continuous Action (COPCAT) datasets). Relative improvements in the training objective from using the proximal point method (left), comparison of test rewards for proximal point vs the simpler gradient-based method (center), and for soft- vs hard-clipping (right). See also Figure 3 for the NoisyMoons dataset.