

# Sample Efficient Policy Search for Optimal Stopping Domains

**Karan Goel**

Carnegie Mellon University  
kgoel93@gmail.com

**Christoph Dann**

Carnegie Mellon University  
cdann@cdann.net

**Emma Brunskill**

Stanford University  
ebrun@cs.stanford.edu

## Abstract

Optimal stopping problems consider the question of deciding when to stop an observation-generating process in order to maximize a return. We examine the problem of simultaneously learning and planning in such domains, when data is collected directly from the environment. We propose GFSE, a simple and flexible model-free policy search method that reuses data for sample efficiency by leveraging problem structure. We bound the sample complexity of our approach to guarantee uniform convergence of policy value estimates, tightening existing PAC bounds to achieve logarithmic dependence on horizon length for our setting. We also examine the benefit of our method against prevalent model-based and model-free approaches on 3 domains taken from diverse fields.

## 1 Introduction

Sequential decision making and learning in unknown environments, commonly modeled as reinforcement learning (RL), is a key aspect of artificial intelligence. An important subclass of RL is optimal stopping processes, where an agent decides at each step whether to continue or terminate a stochastic process. The reward upon termination is a function of the observations seen so far. Many common problems in Computer Science and Operations Research can be modeled within this setting, including the secretary problem [Ferguson, 1989], house selling [Glomer *et al.*, 1998; Lippman and McCall, 1976], American options trading [Jacka, 1991; Mordecki, 2002], product pricing [Feng and Gallego, 1995] and asset replacement [Jiang and Powell, 2015], as well as problems in artificial intelligence like mission monitoring robots [Best *et al.*, 2015], metareasoning about the value of additional computation [Zilberstein, 1995] and automatically deciding when to purchase an airline ticket [Etzioni *et al.*, 2003]. Often the stopping process dynamics are unknown in advance and so finding a good stopping policy (when to halt) requires learning from experience in the environment. As real experience can incur real losses, we desire algorithms that can quickly (with minimal samples) learn good policies that achieve high reward for these problems.

Interestingly, most prior work on optimal stopping has focused on the planning problem: how to compute near-optimal policies given access to the dynamics and reward of the stochastic stopping process [Peskir and Shiryaev, 2006]. Optimal stopping problems can also be framed as a partially observable Markov decision process (POMDP), and there also exists work on learning a good policy for acting in POMDPs, that bounds the number of samples required to identify the near optimal policy out of a class of policies [Kearns *et al.*, 1999; Ng and Jordan, 2000]. However, such work either (i) makes the strong assumption that the algorithm has access to a generative model (ability to simulate from any state) of the stochastic process, which makes this work more suited to improving the efficiency of planning using simulations of the domain, or (ii) can use trajectories directly collected from the environment, but incurs exponential horizon dependence.

In this paper, we consider how to quickly learn a near-optimal policy in a stochastic optimal stopping process with unknown dynamics, given an input class of policies. We assume there is a fixed maximum length horizon for acting, and then make a simple but powerful observation: for stopping problems with process-dependent rewards, the outcomes of a full length trajectory (that is, a trajectory in which the policy only halts after the entire horizon) provide an estimated return for halting after one step, after two steps, and so on, till the horizon. In this way, a single full-length trajectory yields a sample return for any stopping policy. Based on this, we propose an algorithm that first acts by stopping only after the full length horizon for a number of trajectories, and then performs policy search over an input policy class, where the full length trajectories are used to provide estimates of the expected return of each policy considered in the policy class. The policy in the set with the highest expected performance is selected for future use. We provide sample complexity bounds on the number of full length trajectories sufficient to identify a near optimal policy within the input policy class. Our results are similar to more general results for POMDPs [Kearns *et al.*, 1999; Ng and Jordan, 2000], but due to the structure of optimal stopping we achieve two key benefits: our bounds' dependence on the horizon is only logarithmic instead of linear (with a generative model) and exponential (without), and our results apply to learning in stochastic stopping processes, with no generative model required. Simulation results on student tutoring,

ticket purchase, and asset replacement show our approach significantly improves over state-of-the-art approaches.

## 2 Problem Formulation

We consider the standard stochastic discrete-time optimal stopping process setting. As in Tsitsiklis and Van Roy [1999], we assume there is a stochastic process  $P$  that generates observations  $o_1, \dots, o_t$  (they may be vectors). There are two actions: halt or continue the process. The reward model is a known, deterministic function of the sequence of observations and the choice of whether to continue or halt. While there do exist domains where the reward model can be a nondeterministic function of the observations and the actions (such as a medical procedure that reveals the patient’s true condition after a sequence of waiting), most common optimal stopping problems fall within the framework considered here, including the secretary problem (the quality of each secretary is directly observed), house selling (the price for the house from each bidder is known), asset replacement (published guides on the worth of an asset, plus knowledge of the cost of buying a new one), etc. We focus on the episodic setting where there is a fixed maximum time horizon for each process. The finite horizon value of a policy  $V_\pi = E_P[r|\pi]$  is the expected return from following  $\pi$  over a horizon of  $H$  steps, where the expectation is taken over the stochastic process dynamics  $P$ . Note the policy may choose to halt before  $H$  steps. The goal is to maximize return across episodes.

We focus on direct policy search methods (see e.g. [Sutton *et al.*, 1999]). More precisely, we assume as input a parameterized policy class  $\Pi = \{f(\theta) | \theta \in \Theta\}$  where  $\Theta$  is the set of policy parameters. Direct policy search does not require building a model of the domain, and has been very successful in a variety of reinforcement learning (RL) contexts [Deisenroth and Rasmussen, 2011; Levine and Abbeel, 2014].

## 3 Sample Efficient Policy Search

We are particularly interested in domains where evaluation of a policy incurs real cost in the environment, such as stock market options selling. In such settings we wish to find sample efficient methods for doing policy search, that can minimize the number of poor outcomes in the real world. The challenge is that we do not know the stochastic dynamics  $P$  and so it is not possible to, in advance of acting, perform policy search to identify a good policy. Instead we can only obtain information about the domain dynamics by executing policies in the real world. We seek to efficiently leverage such experience to quickly make good decisions.

We now present a simple approach, GFSE (Gather Full, Search and Execute) (Algorithm 1), to do sample efficient policy search. GFSE collects a set of full-length (horizon  $H$ ) trajectories, uses these to evaluate the performance of any policy in the input policy class  $\Pi$ , identifies a good policy, and then executes the resulting policy on all future episodes.

The key insight is in the first step, gathering the data to be used to evaluate the performance of any policy in the policy class  $\Pi$ . Monte Carlo estimation can be used to estimate the expected return of a policy by running it many times. However, this scales poorly with the cardinality of the policy class.

---

### Algorithm 1: Gather Full, Search and Execute (GFSE)

---

Input: policy class  $\Pi$ , search method  $\mathcal{S}$ ,  $\epsilon$ ,  $\delta$   
 $n \leftarrow$  Use Theorem 1 with  $\epsilon$ ,  $\delta$   
 $\Gamma \leftarrow$  Gather  $n$  full trajectories from environment  
 $\pi^* \leftarrow$  Identify policy using  $\mathcal{S}(\Pi, \Gamma)$  \\ \text{evaluation uses } \Gamma

---

Execute  $\pi^*$

---

Building a dynamics model from a set of data is more efficient, as a model can be used to simulate the performance of any policy, but this requires us to make certain assumptions about the domain (for ex. the Markov property) Which can lead to biased estimates. Alternatively, importance sampling can be used to do off-policy evaluation [Precup, 2000] but unfortunately such estimates tend to be very high variance.

However, a simple but powerful observation is that a full-horizon ( $H$ -step) trajectory can be used to yield a sample return for all optimal stopping policies in  $\Pi$ . Given a full length trajectory  $(o_1, \dots, o_H)$ , the performance of a particular policy  $\pi$  can be simulated by providing  $(o_1, o_2, \dots)$  to the target policy until it halts at some time step  $t \leq H$ . Therefore we can take the subsequence of observations  $(o_1, o_2, \dots, o_t)$  and use it to directly compute the return that would have been observed for executing  $\pi$  on this trajectory. A single full-horizon trajectory will provide just one sample of the return of any policy. But a set of  $n$  full-horizon trajectories can be used to provide  $n$  sample returns for a given policy  $\pi$ , thereby providing an empirical estimate of  $V^\pi$ . We can do this off policy evaluation of  $V^\pi$  for any policy in the class  $\Pi$ .

Prior work has shown that given access to a generative model of the domain, policy search can be done in an efficient way by using common random numbers to evaluate policies that act differently in an episode [Kearns *et al.*, 1999; Ng and Jordan, 2000]. In our setting, a full-horizon trajectory is essentially equivalent to having access to a generative model that can produce a single return for any policy. However, access to a full length trajectory can be obtained by running in the environment, whereas generic generative models typically require ”teleportation”: the ability to simulate what would happen next under a particular action given an arbitrary prior history, which is hard unless in a planning scenario in which one already has knowledge of the dynamics process. Our results require weaker assumptions than prior results that use stronger generative models to obtain similar sample efficiency, while also achieving better sample efficiency than approaches with access to similar generative models.

We will shortly provide a sufficient condition on the number of full length trajectories  $n$ , to guarantee that we can evaluate any policy sufficiently accurately to enable policy search to identify a near-optimal policy (within the input policy class). Of course, empirically, we will often wish to select a smaller  $n$ : our simulation experiments will demonstrate that often a small  $n$  still enables us to identify a good policy.

## 4 Theoretical Analysis

We now provide bounds on the sample complexity of GFSE: the number of full length trajectories required to obtain near accurate estimates of all policies in a policy class. This is

sufficient to identify the optimal (or near-optimal) policy in the policy class with the highest expected return.

First, we note that the optimal stopping problems we consider in this paper can be viewed as a particular instance of a POMDP. Briefly, there is some hidden state space, with a dynamics model that determines how the current state transitions to a new state stochastically, given the continue action. The observation is a function of the hidden state, and the reward is also a function of the hidden state and action.

Our main result is that, given a policy class  $\Pi$ , the sample complexity scales logarithmically with the horizon. We make no assumption of access to a generative model. This is a significant improvement over prior sample complexity results for policy search for generic POMDPs and large MDPs [Kearns *et al.*, 1999; Ng and Jordan, 2000] which required access to a generative model of the environment and had a sample complexity that scaled linearly with the horizon. These results can be thought of as bounding the computation/simulation time required during *planning*, when one has access to a generative model that can be used to sample an outcome (reward, observation) given any prior history and action. In contrast, our results apply during *learning*, where the agent has no generative model of the domain, but must instead explore to observe different outcomes. Without a generative model of the domain, sample complexity results for policy search in generic POMDPs when learning scale exponentially with the horizon [Kearns *et al.*, 1999].

Optimal stopping trajectories are related to the trajectory trees of Kearns *et al.* which were used to evaluate the returns of different POMDP policies. For a POMDP with 2 actions, each trajectory tree is a complete binary tree (of depth  $H$ ) rooted at a start state. Nodes in the tree are labeled with a state and observation, and a path from the root to any node in the tree denotes a series of actions taken by a policy. A trajectory tree can be used to evaluate any policy in  $\Pi$ , since every action sequence is part of the tree. However, while for generic POMDPs the size of a trajectory tree is exponential in the horizon, for optimal stopping problems the tree size is linear in the horizon (Figure 1). This allows us to obtain significantly tighter dependence on  $H$  than for generic POMDPs.

Our analysis closely follows the prior sample complexity results of [Kearns *et al.*, 1999]. Kearns *et al.* proceeded by first considering a bound on the VC-dimension of  $\Pi$  when viewed as a set of real-valued mappings from histories to returns. They then use this result to bound the sample complexity needed to get near-accurate estimates of the returns of all policies in the policy class.

We will follow a similar procedure to bound the sample complexity when  $\Pi$  contains a potentially infinite number of deterministic policies.<sup>1</sup> Let  $d = \text{VC}(\Pi)$  be the VC-dimension of our policy class. This is well-defined, since each optimal-stopping policy maps trajectories to 2 actions (binary labeling). Let  $\text{VC}_r(\Pi)$  be the VC-dimension of  $\Pi$  when viewed as a set of real-valued mappings from full trajectories to returns and assume  $V_\pi$  is




Figure 1: The structure of a full trajectory for horizon  $H$ . Each node represents an observation while arrows represent one of the two available actions.

bounded by  $V_{\max}$ . From [Vapnik and Kotz, 1982], we know that  $\text{VC}_r(\Pi)$  can be computed as  $\text{VC}(\mathcal{H})$ , where  $\mathcal{H} = \{I(\pi, \Delta, \cdot) | \pi \in \Pi, \Delta \in [-V_{\max}, V_{\max}]\}$  with  $I(\pi, \Delta, x) = 1$  if  $\pi(x) \geq \Delta$ , and 0 otherwise ( $\pi(x)$  is the return for full trajectory  $x$  under  $\pi$ ).

**Lemma 1.** *Let  $\Pi$  be a set of deterministic optimal-stopping policies with VC-dimension  $d$  when viewed as a set of maps from trajectories to actions. Then, when viewed as a set of maps from the space of all full trajectories to  $[-V_{\max}, V_{\max}]$ ,  $\Pi$  has dimension bounded by,*

$$\text{VC}_r(\Pi) = O(d \log H)$$

*Proof.* Our proof proceeds similarly to Lemma A.1 in Kearns *et al.*. The crucial difference is that our policies operate on a full-trajectory structure that contains  $H$  nodes (Figure 1), rather than Kearns *et al.*'s trajectory trees with  $2^{H+1}$  nodes. In our setting at each point the agent only gets to consider whether to halt or continue, and if the halt action is chosen, the trajectory terminates. This implies that in contrast to standard expectimax trees where the size of the tree depends on the action space as an exponential of the horizon,  $|A|^H$ , in our setting the dependence induced by the actions is only linear in  $H$ . Thus  $\Pi$  can produce a much smaller set of behaviors, and our dependence on  $H$  is logarithmic, rather than polynomial.

More formally, by Sauer's lemma,  $k$  trajectories can be labeled in at most  $(\frac{ek}{d})^d$  ways by  $\Pi$ . First note that  $n$  full trajectories contain at most  $nH$  distinct trajectories across them (one per node; refer to Figure 1 for the structure of full trajectories). Each action labeling of these  $k = nH$  trajectories by  $\Pi$ , corresponds to *selecting*  $n$  paths (1 path per full trajectory), where each path starts at the first observation and ends at a terminal node. The number of possible selections by  $\Pi$  is thus at most  $(\frac{enH}{d})^d$ . Each path can be viewed as mapping a full trajectory to a return, and a selection therefore maps the  $n$  full trajectories to  $n$  real-valued returns.

There are  $nH$  terminal nodes across the  $n$  full trajectories. Thus there are at most  $nH$  distinct real-valued returns on the  $n$  full trajectories under  $\Pi$ . If we set the indicator threshold  $\Delta$  to equal each of these  $nH$  returns in turn, there would be at most  $(\frac{enH}{d})^d$  distinct binary labelings of the full trajectories for each such  $\Delta$ . Thus, the set of indicator functions that define  $\text{VC}_r(\Pi)$  can generate at most  $nH(\frac{enH}{d})^d$  distinct labelings on  $n$  full trajectories. To shatter the  $n$  full trajectories, we set  $nH(\frac{enH}{d})^d \geq 2^n$ , and the result follows.  $\square$

We now proceed similarly to Theorem 3.2 in Kearns *et al.*

**Theorem 1.** *Let  $\Pi$  be a potentially infinite set of deterministic optimal stopping policies and let  $d$  be the VC-dimension of  $\Pi$ . Let  $\Gamma$  be  $n$  full trajectories collected from the environment,*

<sup>1</sup>Similar to Kearns *et al.* (1999) our results extend to finite  $\Pi$  and infinite, stochastic  $\Pi$ , as well as the discounted infinite-horizon case (using an  $\epsilon$ -approximation to  $V_\pi$  with horizon  $H_\epsilon$ ).

and let  $\hat{V}_\pi$  be the value estimates for  $\pi \in \Pi$  using  $\Gamma$ . Let the return  $V_\pi$  be bounded by  $V_{\max}$  for any trajectory. If

$$n = O\left(\left(\frac{V_{\max}}{\epsilon}\right)^2 \left(d \log H + \log \frac{1}{\delta}\right)\right)$$

then with probability at least  $1 - \delta$ ,  $|V_\pi - \hat{V}_\pi| \leq \epsilon$  holds simultaneously for all  $\pi \in \Pi$ .

*Proof.* Let  $X$  be the space of full trajectories. Every policy  $\pi \in \Pi$  is a bounded real-valued map  $\pi : X \rightarrow [-V_{\max}, V_{\max}]$ . Let  $x_1, \dots, x_n \sim P$  be i.i.d. full trajectories generated by the environment dynamics. Using a result of [Vapnik and Kotz, 1982], we have with probability  $1 - \delta$ ,  $\sup_{\pi \in \Pi} |\mathbb{E}_P[\pi(x)] - \frac{1}{n} \sum_{i=1}^n \pi(x_i)| \leq O\left(V_{\max} \sqrt{\frac{VC_r(\Pi) \log \frac{n}{VC_r(\Pi)} + \log \frac{1}{\delta}}{n}}\right)$ . Substitute  $\mathbb{E}_P[\pi(x)] = V_\pi$ ,  $\frac{1}{n} \sum_{i=1}^n \pi(x_i) = \hat{V}_\pi$ ,  $VC_r(\Pi) = O(d \log H)$  in the inequality and upper-bound by  $\epsilon$  to get the result.  $\square$

In practice it may be impossible for us to evaluate every policy in  $\Pi$ , and then select the one with the best estimated mean. In such cases, we can use a different search method ( $\mathcal{S}$  in Algorithm 1) to find a local optima in  $\Pi$ , while using our bound to ensure that policy values are estimated accurately.

Lastly, we discuss [Tsitsiklis and Van Roy, 1999], who estimate  $Q$  values for finite-horizon Markov optimal stopping problems using a linear combination of basis functions, and then use that to find a threshold policy. They outline a procedure to tune the basis function weights that asymptotically guarantees their policy value's convergence to the best basis-function approximation. Under their assumptions, if we construct a policy class using basis functions, we inherit the useful convergence results relying on their search procedure, along with retaining our finite sample complexity results.

## 5 Experiments

We now demonstrate the setting we consider is sufficiently general to capture several problems of interest and that our approach, GFSE, can improve performance in optimal stopping problems over some state-of-the-art baselines.

### Ticket Purchase

Many purchasing problems can be posed as an optimal stopping process where the return from stopping is simply the advertised cost. We consider deciding when to purchase an airline ticket for a later trip date in order to minimize cost. The opaque way in which prices are set, and competitive pricing makes this domain difficult to model. Prior work [Etzioni *et al.*, 2003; Groves and Gini, 2015] has focused on identifying features to create sophisticated models that make good purchase decisions. Surprisingly, it can be hard to improve on an earliest purchase baseline that buys after the first observation.

We use data from Groves and Gini (2015) who collected real pricing data for a fixed set of routes over a period of 2 years, querying travel sites regularly to collect price information. Each route has several departure dates distributed over the 2 year period. For a price observation sequence of length

Table 1: Mean expenditure of deploying different policies on the test set for ticket purchase. Earliest purchase buys immediately, latest purchase waits till the departure date.

Method	NYC-MSP	MSP-NYC	SEA-IAD
Ours ( $\Pi_{\text{simple}}$ )	\$355	\$374	\$565
Ours ( $\Pi_{\text{complex}}$ )	<b>\$351</b>	<b>\$344</b>	<b>\$560</b>
Earliest purchase	\$380	\$383	\$578
Latest purchase	\$631	\$647	\$938
Best possible price	\$307	\$306	\$513

$T$ , a customer could commence his ticket search at any point in the sequence (e.g. some customer starts 60 days before departure while another only a week before). Thus, we consider all such commencement points separately to get  $T$  distinct full trajectories (similar to [Groves and Gini, 2015]).

We construct a parameterized policy class ( $\Pi_{\text{simple}}$ ) based on Ripper's decision rules in [Etzioni *et al.*, 2003]: WAIT if (curr.price >  $\theta_0$  AND days.to.depart >  $\theta_1$ ) else BUY, where BUY corresponds to halting. We also constructed a more complex class ( $\Pi_{\text{complex}}$ ) with 6 parameters, that learns different price thresholds depending on how far the departure date is. We consider nonstop flights on 3 routes, NYC-MSP, MSP-NYC and SEA-IAD, training/testing each separately.

Our method, GFSE collects full length trajectories during the first 200 days ( $\sim 1000$  trajectories) and uses them to construct a single stopping policy. It performs a simple policy search by sampling and evaluating 500 policies randomly from the policy space. It then uses the best identified policy to simulate ticket purchasing decisions for departure dates occurring during the remaining part of the 2 years ( $\sim 2000$  trajectories). We restrict the data to departure dates that contain at-least 30 price observations.<sup>2</sup>

Results on the test sets are shown in Table 1. Our policy search method succeeds in finding a policy that leads to non-trivial improvement over the difficult earliest purchase baseline. Our improvements are in line with prior approaches specifically designed for this particular domain.<sup>3</sup> These results highlight how our setting can capture important purchasing tasks and how our approach, even with a simple policy search, can find policies with significantly better performance than competitive domain-specific baselines.

### Tutoring and Asset Replacement

We now consider 2 simulated domains and compare GFSE to several approaches for learning to act quickly in these domains. Unless specified, all results are averaged over 20 rounds and error bars indicate 95% confidence intervals.

**Baselines.** One natural idea is to proceed as GFSE, but use the gathered data to build parametric domain models that can be used to estimate the performance of potential policies. We

<sup>2</sup>We found that shorter trajectories were collected close to the departure date, where prices fluctuate more and for which our illustrative policy classes are inadequate. In such cases, our method adopted a risk-averse earliest purchase policy.

<sup>3</sup>Unfortunately, the authors were unable to provide us with the train/test split used in [Groves and Gini, 2015].

call these "model-based" approaches. A second idea is to consider the initial set of collected data as a budget of free exploration, and instead use this budget to do Monte Carlo on-policy evaluation of a set of policies.

Of course, doing all exploration, as we do in GFSE, is not always optimal. We also consider a state-of-the-art approach for quickly identifying the global optima of a function where the function is initially unknown and each function evaluation is expensive, Bayesian Optimization (BO). Multiple papers have shown BO can be used to speed online policy search for reinforcement learning tasks [Wilson *et al.*, 2014; Deisenroth and Rasmussen, 2011]. Given the policy class, BO selects a policy to evaluate at each step, and maintains estimates over the expected value of every policy. We use Yelp’s MOE for BO [Yelp, 2016] with a Gaussian kernel and the popular expected improvement heuristic for picking policies from  $\Pi$ . The hyper-parameters for BO are picked by a separate optimization to find maximum-likelihood estimates.

**Simulated Student Learning.** We first consider a simulated student tutor domain. A number of tutoring systems use mastery teaching, in which a student is provided with practice examples until they are estimated to have mastered the material. This is an optimal stopping problem because at each time step, after observing whether a student got the activity correct or not, the tutor can decide whether to halt or continue providing the student with additional practice. On halting, the student is given the next problem in the sequence; the objective is to maximize the score on this 'posttest', while giving as few problems as possible overall. As is popular in the literature, we model student learning using the Bayesian Knowledge Tracing (BKT) model [Corbett and Anderson, 1995]. BKT is a 2-state Hidden Markov Model (HMM) with the state capturing whether the student has mastered the skill or not. Within the HMM, 4 probabilities –  $p_i$  (prior mastery),  $p_t$  (transition to mastery),  $p_g$  (guess) and  $p_s$  (slip) describe the model. To simulate student data, we fix BKT parameters<sup>4</sup>  $p_i = 0.18, p_t = 0.2, p_g = 0.2, p_s = 0.1$  and generate student trajectories using this BKT model for  $H = 20$  problems.

For GFSE, we consider two policy classes, both of which halt when the probability of the student’s next response (according to the model in use) being correct crosses a threshold. Thus, we halt if  $\Pr(o_t = \text{correct} | o_1, \dots, o_{t-1}) > \theta_0$  where  $\theta_0$  is some threshold. In fact, policies of this kind are widely used in commercial tutoring systems [Koedinger *et al.*, 2013]. If we use the BKT model to implement this policy class, it is parameterized by  $(p_i, p_t, p_g, p_s, \theta_0)$ .  $\Pi$  then contains all possible instantiations of these parameters for our model-free approach to search over. We also consider a policy class based on another popular educational data mining model of student learning: Additive Factors Model (AFM) [Draney *et al.*, 1995]. AFM is a logistic regression model used to predict the probability that a student will get the next problem correct given their past responses. Thus,  $\Pr(o_t = \text{correct} | o_1, \dots, o_{t-1}) = \frac{1}{1 + e^{-(\beta_1 + \beta_2 n_c)}}$  where  $n_c$  is the number of correct past attempts.

<sup>4</sup>Our results hold for other instantiations of these parameters as well. See [Ritter *et al.*, 2009] for other reasonable parameter settings.




Figure 2: Comparison of best policy found on varying budget under (a) matched model setting; (b) model mismatch setting.

We first note that GFSE is significantly more effective than taking the same budget of exploration, and using it to evaluate each policy in an on-policy manner using Monte Carlo (MC) estimation. More precisely, we sample  $k = 100$  policies from the BKT policy class, and fix a budget of  $B \in \{100, 1000\}$  trajectories. GFSE uses  $B$  trajectories to evaluate all  $k$  policies while MC runs every policy on  $\frac{B}{k}$  trajectories (e.g. 1 trajectory/policy for  $B = 100$ ) and selects the one with the highest mean performance. Averaging results across 20 separate runs, we found that GFSE identifies a much better policy; MC chose poor policies because it is misled by the potential performance of a policy due to the limited data.

We also explored the performance of building a model of the domain, both in the setting when the model matches the true domain (here, a BKT model), and a model-mismatch case, where the policy class is based on a student AFM model (which does not match the BKT process dynamics). We use maximum likelihood estimation to fit the assumed model’s parameters given the collected data and then separately optimize over the threshold parameters  $\theta_0$ . We compare, on varying the budget  $B$ : (a) GFSE; (b) model-based; (c) BO. All results are averaged over 50 trials.

The results of this experiment are shown in Figure 2. Our approach does well in both settings, quickly finding a near optimal policy. As one would expect, the model-based approach does well under the matched model setting, making full use of the knowledge of the underlying process dynamics. However, on fitting the mismatched AFM model, the model-based approach suffers. As has been noted by prior work [Mandel *et al.*, 2014] model-fitting procedures focus on maximizing the likelihood of the observed data rather than trying to directly identify a policy that is expected to perform well. BO can find a good policy, but takes more samples to do so.

Since BO is an online approach whereas GFSE uses a fixed budget of exploration, we also compare the averaged cumulative performance of BO to variants of GFSE in Figure 3. This mimics a scenario where we care about online performance on every individual trajectory, rather than having access to a fixed budget before deploying a policy. For our method, we can choose to collect more or less full trajectories before finding the best policy. Interestingly, if we use 5 trajectories as the initial budget to collect full length trajectories, GFSE meets or exceeds BO performance in this setting in both the matched and mismatched model cases, within 10 trajectories.

BO suffers from the highly stochastic returns of policies in this setting. For more efficient data reuse, we also consider a




Figure 3: Average cumulative performance on the simulated student domain. (a) matched model; (b) model-mismatch. GFSE- $k$  collects  $k$  initial full-length trajectories before identifying and then executing the best policy.




Figure 4: Cumulative performance of augmented methods on (a) matched model; (b) model-mismatch.

variant of BO (BO-REuse) where we evaluate each proposed policy online and also using previously collected trajectories, yielding a more robust estimate of the policy’s performance. Similarly, for GFSE we deploy a policy using an initial budget of trajectories, and then use its on-policy trajectory (in addition to earlier trajectories) to rerun policy search and identify another policy for the next time step (GFSE-RE). Figure 4 shows this improved both methods (especially in the mismatch case), with our approaches still performing best.

**Asset Replacement.** Another natural problem that falls into our optimal stopping problem is when to replace a depreciating asset (such as a car, machine, etc). For simulation, we use a model described in [Jiang and Powell, 2015]. Variants of this model are widely used in that field [Feldstein and Rothschild, 1974; Rust, 1987]. In the model, observations are  $d$  dimensional vectors of the form  $(X, Y_1, \dots, Y_{d-1})$ . Each asset starts at a fixed valuation  $X = X_{\max}$  which depreciates stochastically<sup>5</sup> while emitting observations  $Y_1, \dots, Y_{d-1}$  at every time step. The reward function used incorporates the cost of replacement (which increases over time), the utility derived from the asset and a penalty if the asset becomes worthless before replacement. We use  $d = 3$  for experiments.

We construct a logistic threshold policy class; replacing the asset if  $\frac{1}{1+e^{-(\beta_1+\beta_2 \cdot \text{depr})}} > \beta_3$  where  $\text{depr}$  is the total depreciation from  $X_{\max}$  seen so far (normalized to lie in  $[0, 1]$ ). In addition to the approaches seen before, we also include baseline policies that choose to (i) replace the asset immediately; (ii) never replace. Lastly, we include the optimal value (known only in hindsight) for reference.

The results are shown in Fig 5. Surprisingly, our method




Figure 5: Results from asset replacement.

outperforms competing methods by a considerable margin. It appears that our chosen policy class is tricky to optimize over: most policies in the space perform poorly. For 500 random policies chosen from this space, the mean cost is around 240 with a 95% confidence interval of only 16. However, the domain itself is not very noisy, with robust value estimation requiring less than 5 trajectories (see Figure 5). This enables our method to consistently find a good policy even with a low budget: one that corresponds to replacing the asset when depreciation is around 50%. BO improves slowly; either sampling bad policies due to the sparse nature of the space, or disbelieving the estimate of a good policy due to the bad policies surrounding it. Manually adjusting the BO hyperparameters to account for this did not improve performance significantly.

## 6 Discussion and Conclusion

GFSE performed well, outperforming state-of-the-art algorithms and common baselines, in a variety of simulations of important domains. While we randomly searched over policies in relatively simple policy classes for illustration, more sophisticated search methods and policy classes could be employed as well. One extension is in using shorter trajectories that terminate before the horizon for policy evaluation (similar to how full trajectories are used). This is useful in a scenario where we get trajectories on-policy using the best policy found by GFSE. We can then rerun policy search with all trajectories (full length or short) collected so far. Our policy value estimates will be biased in this case, since only a policy that halts earlier than a shorter trajectory can use it for evaluation. Values for policies that halt later may be overestimated (higher variance of estimation due to fewer trajectories), pushing us to pick them. If the number of evaluations per policy exceeds the number in Theorem 1, our estimates would remain within  $\epsilon$  of the true values (with high probability), which would minimize the effect of this bias. As we saw from Figure 4, this (GFSE-RE) works well empirically.

To summarize, we introduced a method for learning to act in optimal stopping problems, which reuses full length trajectories to perform policy search. Our theoretical analysis and empirical simulations demonstrate that this simple observation can lead to benefits in sample complexity and practice.

## Acknowledgments

We appreciate the financial support of a NSF BigData award #1546510, a Google research award and a Yahoo gift.

<sup>5</sup>Details of the model can be found in [Jiang and Powell, 2015].

## References

- [Best *et al.*, 2015] Graeme Best, Wolfram Martens, and Robert Fitch. A spatiotemporal optimal stopping problem for mission monitoring with stationary viewpoints. In *Robotics: Science and Systems*, 2015.
- [Corbett and Anderson, 1995] Albert T. Corbett and John R. Anderson. Knowledge tracing: Modelling the acquisition of procedural knowledge. *User Model. User-Adapt. Interact.*, 4(4):253–278, 1995.
- [Deisenroth and Rasmussen, 2011] Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *ICML*, pages 465–472, 2011.
- [Draney *et al.*, 1995] Karen L Draney, Peter Pirolli, and Mark Wilson. A measurement model for a complex cognitive skill. *Cognitively diagnostic assessment*, pages 103–125, 1995.
- [Etzioni *et al.*, 2003] Oren Etzioni, Rattapoom Tuchinda, Craig A. Knoblock, and Alexander Yates. To buy or not to buy: mining airfare data to minimize ticket purchase price. In *KDD*, pages 119–128, 2003.
- [Feldstein and Rothschild, 1974] Martin S Feldstein and Michael Rothschild. Towards an economic theory of replacement investment. *Econometrica: Journal of the Econometric Society*, pages 393–423, 1974.
- [Feng and Gallego, 1995] Youyi Feng and Guillermo Gallego. Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science*, 41(8):1371–1391, 1995.
- [Ferguson, 1989] Thomas S Ferguson. Who solved the secretary problem? *Statistical science*, pages 282–289, 1989.
- [Glomer *et al.*, 1998] Michel Glomer, Donald R Haurin, and Patric H Hendershott. Selling time and selling price: The influence of seller motivation. *Real estate economics*, 26(4):719–740, 1998.
- [Groves and Gini, 2015] William Groves and Maria L. Gini. On optimizing airline ticket purchase timing. *ACM TIST*, 7(1):3:1–3:28, 2015.
- [Jacka, 1991] SD 1 Jacka. Optimal stopping and the american put. *Mathematical Finance*, 1(2):1–14, 1991.
- [Jiang and Powell, 2015] Daniel R. Jiang and Warren B. Powell. An approximate dynamic programming algorithm for monotone value functions. *Operations Research*, 63(6):1489–1511, 2015.
- [Kearns *et al.*, 1999] Michael J. Kearns, Yishay Mansour, and Andrew Y. Ng. Approximate planning in large pomdps via reusable trajectories. In *NIPS*, pages 1001–1007, 1999.
- [Koedinger *et al.*, 2013] Kenneth R Koedinger, Emma Brunskill, Ryan SJD Baker, Elizabeth A McLaughlin, and John Stamper. New potentials for data-driven intelligent tutoring system development and optimization. *AI Magazine*, 34(3):27–41, 2013.
- [Levine and Abbeel, 2014] Sergey Levine and Pieter Abbeel. Learning neural network policies with guided policy search under unknown dynamics. In *NIPS*, pages 1071–1079, 2014.
- [Lippman and McCall, 1976] Steven A Lippman and John McCall. The economics of job search: A survey. *Economic inquiry*, 14(2):155–189, 1976.
- [Mandel *et al.*, 2014] Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. Offline policy evaluation across representations with applications to educational games. In *AAMAS*, pages 1077–1084. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [Mordecki, 2002] Ernesto Mordecki. Optimal stopping and perpetual options for lévy processes. *Finance and Stochastics*, 6(4):473–493, 2002.
- [Ng and Jordan, 2000] Andrew Y. Ng and Michael I. Jordan. PEGASUS: A policy search method for large mdps and pomdps. In *UAI*, 2000.
- [Peskir and Shiryaev, 2006] Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006.
- [Precup, 2000] Doina Precup. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80, 2000.
- [Ritter *et al.*, 2009] Steven Ritter, Thomas K. Harris, Tristan Nixon, Daniel Dickison, R. Charles Murray, and Brendon Towle. Reducing the knowledge tracing space. In *Educational Data Mining*, pages 151–160, 2009.
- [Rust, 1987] John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033, 1987.
- [Sutton *et al.*, 1999] Richard S. Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, pages 1057–1063, 1999.
- [Tsitsiklis and Van Roy, 1999] John N Tsitsiklis and Benjamin Van Roy. Optimal stopping of markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control*, 44(10):1840–1851, 1999.
- [Vapnik and Kotz, 1982] Vladimir Naumovich Vapnik and Samuel Kotz. *Estimation of dependences based on empirical data*, volume 40. Springer-Verlag New York, 1982.
- [Wilson *et al.*, 2014] Aaron Wilson, Alan Fern, and Prasad Tadepalli. Using trajectory data to improve bayesian optimization for reinforcement learning. *Journal of Machine Learning Research*, 15(1):253–282, 2014.
- [Yelp, 2016] Yelp. Metric optimization engine. <https://github.com/Yelp/MOE>, 2016.
- [Zilberstein, 1995] Shlomo Zilberstein. Operational rationality through compilation of anytime algorithms. *AI Magazine*, 16(2):79, 1995.