

Model Selection for Offline Reinforcement Learning: Practical Considerations for Healthcare Settings

Shengpu Tang

*Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI, USA*

TANGSP@UMICH.EDU

Jenna Wiens

*Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI, USA*

WIENSJ@UMICH.EDU

Abstract

Reinforcement learning (RL) can be used to learn treatment policies and aid decision making in healthcare. However, given the need for generalization over complex state/action spaces, the incorporation of function approximators (e.g., deep neural networks) requires model selection to reduce overfitting and improve policy performance at deployment. Yet a standard validation pipeline for model selection requires running a learned policy in the actual environment, which is often infeasible in a healthcare setting. In this work, we investigate a model selection pipeline for offline RL that relies on off-policy evaluation (OPE) as a proxy for validation performance. We present an in-depth analysis of popular OPE methods, highlighting the additional hyperparameters and computational requirements (fitting/inference of auxiliary models) when used to rank a set of candidate policies. We compare the utility of different OPE methods as part of the model selection pipeline in the context of learning to treat patients with sepsis. Among all the OPE methods we considered, fitted Q evaluation (FQE) consistently leads to the best validation ranking, but at a high computational cost. To balance this trade-off between accuracy of ranking and computational efficiency, we propose a simple two-stage approach to accelerate model selection by avoiding potentially unnecessary computation. Our work serves as a practical guide for offline RL model selection and can help RL practitioners select policies using real-world datasets. To facilitate reproducibility and future extensions, the code accompanying this paper is available online.¹

1. Introduction

In recent years, researchers have begun to explore the potential of reinforcement learning (RL) to aid in sequential decision making in the context of clinical care (Yu et al., 2019). For example, Prasad et al. (2017) considered the problem of weaning from mechanical ventilation, Komorowski et al. (2018) used RL to learn optimal treatments of sepsis, and Futoma et al. (2020a,b) explored hypotension management in the ICU. This is an important step towards realizing the promise of RL beyond game-like domains. However, many challenges remain. To date, the most successful applications of RL have been in settings that assume an online paradigm in which the agent can interact with the environment (Mnih et al., 2015; Silver et al., 2017; Li et al., 2010; Kalashnikov et al., 2018), an assumption that often does

1. https://github.com/MLD3/OfflineRL_ModelSelection

not hold in other contexts. In healthcare settings, one typically only has access to historical data that were collected “offline”, thus necessitating offline RL techniques (Levine et al., 2020). In addition, the state space associated with clinical decision making is often high-dimensional and/or continuous. Capturing such complex relationships may require learning parameterized functions (e.g., neural networks) (Mnih et al., 2015). Thus, model selection, i.e., selecting the appropriate hypothesis class and hyperparameters of the function approximators for learning policies, is crucial in offline RL to prevent overfitting. Despite the importance of model selection, we currently lack a well-accepted training-validation framework for offline RL.

One cannot simply rely on the training-validation framework used in supervised learning, in which the final model is selected based on validation performance. In RL, measuring the equivalent of “validation performance” requires running the learned policy in the actual environment. While this is common practice when accurate simulators are available (Fujimoto et al., 2019a,b; Kumar et al., 2019; Larocche et al., 2019; Liu et al., 2020), it is not always feasible in healthcare settings; it may be costly or even dangerous to try out multiple untested treatment policies on real patients in order to find out which one is the best. In healthcare, researchers must rely on off-policy evaluation (OPE) techniques for model selection in offline RL. However, this approach has its own challenges. Since many OPE estimators have been developed and they can only approximate validation performance, it is unclear how useful different OPE methods are for model selection. Paine et al. (2020) and Fu et al. (2021) found FQE to be a useful strategy for hyperparameter selection in offline RL on several benchmark control tasks, yet the practical limitations and tradeoffs of these OPE methods have not been systematically examined. Furthermore, OPE methods often require additional model selection, i.e., they involve *additional* hyperparameters that must be set *a priori*. In addition, nearly every OPE method requires auxiliary models that must also be learned from data, leading to (sometimes substantial) increases in computational costs. The lack of guidelines for offline RL model selection has led the majority of empirical RL analyses on healthcare data to rely on pre-specified hyperparameters during training (Gottesman et al., 2018; Tang et al., 2020), but this does not necessarily result in learning the best policy nor does it enable fair comparisons between different RL algorithms.

In this work, we propose a training-validation framework that relies on OPE for model selection in offline RL settings. We investigate four popular OPE methods, highlighting practical considerations pertaining to their dependence on additional hyperparameters, auxiliary models, and their associated computational costs when used to rank a set of candidate policies. To quantify the effectiveness of each OPE method for model selection, we perform empirical experiments using a sepsis treatment task (Oberst and Sontag, 2019; Futoma et al., 2020a). We consider common model selection scenarios including early stopping and/or choosing neural network architectures, for both discrete state spaces (tabular) and continuous state spaces (function approximation). Given the relative computational costs and ranking quality of different OPE estimators, we propose a simple two-stage approach that combines two OPE estimators to avoid expensive computation on low-quality policies. First, we use a more efficient but less accurate OPE method to identify a promising subset of candidate policies, and second, we use a more accurate but less efficient OPE method to select a final policy from the pruned subset.

Generalizable Insights about Machine Learning in the Context of Healthcare

We tackle the problem of model selection for offline RL applied to observational datasets, such as those in healthcare applications. Selecting appropriate hyperparameters for model-free value-based RL is crucial to prevent overfitting. Our contributions can be summarized as follows:

- We analyze four popular OPE methods in the context of estimating validation performance of learned policies, highlighting their dependence on additional hyperparameters, auxiliary models, and their computational requirements.
- On the problem of using RL to learn policies for treating patients with sepsis, we empirically evaluate the effectiveness of using OPE validation performance for model selection, addressing practical considerations.
- We propose a simple two-stage selection approach that effectively combines OPE estimators, balancing computational efficiency with policy performance.

RL shows promise in discovering better treatment policies from historical data, but the challenges of learning and evaluating policies in offline settings have impeded its wider adoption in high-stakes clinical domains. Our work provides one of the first empirical demonstrations of offline RL model selection using OPE methods with a focus on practical considerations; it sheds light on the pros and cons of various OPE methods for this purpose and can inform future research and enable fairer comparisons when applying offline RL to problems in healthcare. Though inspired by recent applications of RL to healthcare, many of the insights here generalize beyond healthcare decision making to other high-stakes domains where accurate simulators are unavailable, e.g., intelligent tutoring systems, dialogue systems, and online advertisement delivery.

2. Offline Reinforcement Learning - Problem Setup

We consider Markov decision processes (MDPs) defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p, r, \mu_0, \gamma)$, where \mathcal{S} and \mathcal{A} represent the state and action spaces, $p(s'|s, a)$ and $r(s, a)$ specify the transition and reward models, $\mu_0(s)$ is the initial state distribution, and $\gamma \in [0, 1]$ is the discount factor. In this work, we consider settings with discrete/continuous state spaces and a discrete action space. A probabilistic policy $\pi(a|s)$ specifies a mapping from each state to a probability distribution over actions. When the policy is deterministic, $\pi(s)$ refers to the action with $\pi(a|s) = 1$. In RL, we aim to find an optimal policy $\pi^* = \arg \max_{\pi} J(\pi; \mathcal{M})$ that has the maximum expected performance in \mathcal{M} . The value of policy π with respect to \mathcal{M} is defined as $v(\pi) = J(\pi; \mathcal{M}) = \mathbb{E}_{s \sim \mu_0}[V^\pi(s)]$; here, the state-value function is defined as $V^\pi(s) = \mathbb{E}_{\pi} \mathbb{E}_{\mathcal{M}} [\sum_{t=1}^{\infty} \gamma^{t-1} r_t \mid s_1 = s]$. The action-value function, $Q^\pi(s, a)$, is defined by further restricting the action taken from the starting state.

Throughout this paper, we focus on the *offline* setting. In contrast to an *online* setting in which one can interact with \mathcal{M} (either directly or via a simulator), in offline RL, one only has access to data that were previously collected. More precisely, one has access to an observational dataset $\mathcal{D}_{\text{OBS}} = \{s_i, a_i, r_i, s'_i\}_{i=1}^N$ consisting of N transitions, collected from \mathcal{M} by one or multiple behavior policies denoted by π_b . The dataset can also be considered as a collection of m episodes, $\mathcal{D}_{\text{OBS}} = \{\tau_j\}_{j=1}^m$, where each episode is $\tau = (s_1, a_1, r_1, \dots, s_L, a_L, r_L)$ with

L denoting the length of episode τ . Given the challenges such a setting presents (i.e., limited exploration), we cannot necessarily identify an optimal policy (Lange et al., 2012; Levine et al., 2020); instead, we aim to learn a policy with low suboptimality $J(\pi^*; \mathcal{M}) - J(\pi; \mathcal{M})$. We focus on model-free, value-based methods for offline policy learning – specifically, fitted Q iteration (FQI) (Ernst et al., 2005; Riedmiller, 2005; see Alg. A.1 in Appendix A) – as it applies to problems with either discrete state spaces (the tabular setting) or continuous state spaces (the function approximation setting, where states are represented as feature vectors $\mathbf{x}(s) \in \mathbb{R}^d$).

Given a policy π , its performance (i.e., $J(\pi; \mathcal{M})$), the expected cumulative reward in \mathcal{M}) can be estimated from historical data collected by some other policy (or policies) π_b using off-policy evaluation (OPE) (Voloshin et al., 2019). OPE is inherently difficult because it requires estimating counterfactuals: we want to understand what *would have* happened had the agent acted differently from how it acted in historical data. We discuss four popular OPE methods in Section 3.2 and show how they can be used to estimate validation performance in the context of model selection for offline RL.

3. Offline RL Model Selection using OPE

We present a model selection framework for offline RL based on OPE. First, we analyze four common OPE methods in the context of model selection. We highlight practical considerations including their dependence on additional hyperparameters and auxiliary models, as well as computational requirements. Second, we propose a simple two-stage approach for model selection that combines two OPE estimators in order to balance the trade-off between ranking quality and computational efficiency.

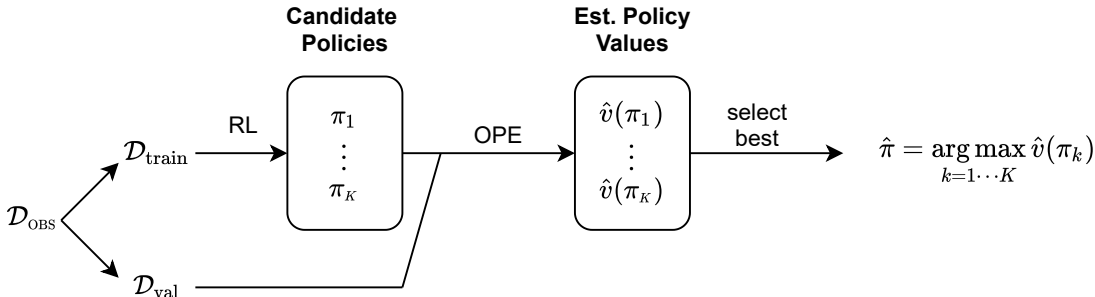


Figure 1: The observational dataset \mathcal{D}_{OBS} is first split into two mutually exclusive parts, $\mathcal{D}_{\text{train}}$ and \mathcal{D}_{val} . The training data $\mathcal{D}_{\text{train}}$ are used to derive a set of candidate policies by running some RL algorithm (e.g., FQI). Each policy is evaluated using an OPE method on the validation data \mathcal{D}_{val} to obtain its estimated policy value $\hat{v}(\pi_k)$. Finally, we select the policy that achieves the highest estimated validation performance.

3.1. A Model Selection Framework for Offline RL

The model selection framework commonly used in supervised learning (Murphy, 2012; Goodfellow et al., 2016) can be adapted for offline RL, where OPE methods are used to estimate the validation performance (Figure 1). Given an observational dataset, we first split it into two mutually exclusive parts: a training set $\mathcal{D}_{\text{train}}$ and a validation set \mathcal{D}_{val} , assumed to follow the same data distribution (the same MDP and the same behavior(s)). During training, a set of K candidate policies $\{\pi_k\}_{k=1}^K$ are derived from the $\mathcal{D}_{\text{train}}$ using an algorithm such as FQI (Alg. A.1), where each policy is learned with a different hyperparameter setting (e.g., training iterations, neural network architectures). Using OPE, we then assign a scalar score to each candidate policy π_k by estimating its value using \mathcal{D}_{val} . The scores collectively result in a ranking that can then be used to identify the best policy. While seemingly straightforward, such a model selection framework has not been consistently applied, in part because of the practical challenges that one encounters in offline settings.

3.2. Using OPE for Model Selection

Compared to measuring validation performance in a supervised learning setting (by checking predictions against ground-truth labels), measuring validation performance in offline RL cannot be simply done using the observed returns from the validation set. This is because from the same state, the evaluation policy may choose a different action (or a different distribution over actions) compared to the behavior policy. Thus, we have to use OPE approaches, which themselves rely on additional hyperparameters or learning additional models. In this section, we provide a qualitative comparison of popular OPE methods. We describe each OPE and its mathematical formalization. In addition, we describe any additional hyperparameters that must be set *a priori*, dependencies on auxiliary models, and computational requirements when used for model selection. As shown in Figure 1, we use the validation data \mathcal{D}_{val} as input to OPE, but to simplify the exposition below, we use \mathcal{D} for notational convenience. We assume \mathcal{D} consists of m episodes and N transitions, and there are K candidate policies. For computational complexity, we focus on the function approximation setting that involves fitting classification/regression models; the equivalent of these model dependencies in the tabular setting can be obtained efficiently via maximum likelihood estimation (see Appendix B).

3.2.1. WEIGHTED IMPORTANCE SAMPLING (WIS)

Due to differences between the evaluation policy and the behavior policy, each episode can be more (or less) likely to occur under the evaluation policy compared to the behavior policy. Importance sampling estimates the value of the evaluation policy from the validation data by re-weighting episodes according to how relatively likely they are to occur (Păduraru et al., 2013; Voloshin et al., 2019).

Formalization. Given a policy π that we aim to evaluate and a behavior policy π_b that generated episode $\tau = (s_1, a_1, r_1, \dots, s_L, a_L, r_L)$, we define the per-step importance ratio $\rho_t = \frac{\pi(a_t|s_t)}{\pi_b(a_t|s_t)}$ and the cumulative importance ratio $\rho_{1:t} = \prod_{t'=1}^t \rho_{t'}$. The WIS estimator is calculated as

$$\hat{v}_{\text{WIS}}(\pi) = \frac{1}{m} \sum_{j=1}^m G_{\text{WIS}}(\tau_j; \pi), \text{ where } G_{\text{WIS}}(\tau_j; \pi) = \frac{\rho_{1:L(j)}^{(j)}}{w_{L(j)}} \left(\sum_{t=1}^{L(j)} \gamma^{t-1} r_t^{(j)} \right),$$

where $w_t = \frac{1}{m} \sum_{j=1}^m \rho_{1:t}^{(j)}$ is the average cumulative importance ratio at horizon t . We use WIS (a biased but consistent estimator) rather than ordinary importance sampling (an unbiased estimator) because it can provide more stable performance (Păduraru et al., 2013; Voloshin et al., 2019).

Hyperparameters. When the evaluation policy π is deterministic such that $\pi(a|s)$ is exactly 1 for a single action and 0 for all others, an episode will receive zero weight if π and π_b take different actions at *any* timestep. Such episodes are thus effectively discarded, and only episodes with non-zero weights are included in WIS calculation (Gottesman et al., 2018). This can be especially problematic with small sample sizes and long episodes, or when π and π_b differ significantly. To overcome this issue, it is common to *soften* policy π with the hope of increasing the effective sample size and reducing variance (Raghu et al., 2017a,b; Komorowski et al., 2018): e.g., $\tilde{\pi}(a|s) = \mathbb{1}_{a=\pi(s)}(1-\varepsilon) + \mathbb{1}_{a \neq \pi(s)}(\frac{\varepsilon}{|\mathcal{A}|-1})$. While this results in a value estimate for $\tilde{\pi}$ rather than for π , it is usually sufficiently close. The hyperparameter ε is typically set to a small value using rules of thumb (e.g., $\varepsilon = 0.01$).

Dependencies on Auxiliary Models. Since we typically do not know the true behavior policy, we need to estimate it from data (Hanna et al., 2019; Chen et al., 2019). This can be obtained as either $\hat{\pi}_b(a|s) = \frac{\text{count}(s,a)}{\text{count}(s)}$ for discrete states, or for continuous states, by fitting a multiclass classification model that predicts the action distribution of the behavior policy given a state, trained using the softmax cross-entropy loss (Torabi et al., 2018).

Computational Cost. WIS requires fitting a single model $\hat{\pi}_b$ over all observed state-action pairs $\{\mathbf{x}(s_i), a_i\}_{i=1}^N$ and performing one inference run on all observed states $\{\mathbf{x}(s_i)\}_{i=1}^N$, resulting in $O(N)$ complexity for both fitting and inference. Importantly, the computational cost can be considered amortized since it does not depend on K , the number of candidate policies.² This allows us to apply WIS to a large candidate set with relatively little increase in computation.

3.2.2. APPROXIMATE MODEL (AM)

In the online setting, we evaluate a policy by running it in the actual environment. While we do not have access to the true environment in the offline case, we can build an approximate model of the environment from \mathcal{D} . AM is a model-based approach that estimates the policy value by first learning a model of the environment and then performing policy evaluation using the learned model (Jiang and Li, 2016; Voloshin et al., 2019). As the name suggests, AM is only an approximation but may be sufficient for providing a reasonable estimate of the policy value.

2. To be precise, applying WIS on each policy will require new importance ratios to be computed, but the cost of such arithmetic operations is minimal compared to performing fitting/inference of a parameterized model (e.g., neural networks).

Formalization. To learn a set of models for AM, we need to fit the transition dynamics $\hat{p}(s'|s, a)$ and the reward function $\hat{r}(s, a)$, and if necessary, the initial state distribution $\hat{\mu}_0(s)$, and/or the termination model $\hat{p}_{\text{term}}(s, a)$ for episodic problems. Given these estimated parameters of the environment, we can estimate the value functions \hat{V}_{AM} and/or \hat{Q}_{AM} by performing analytic or iterative policy evaluation for the given policy (only applicable to the tabular setting, i.e., discrete state and action spaces), or use the learned model as a simulator to obtain average returns over Monte-Carlo rollouts (Jiang and Li, 2016; Voloshin et al., 2019). The final policy value can be calculated as the average value of initial states (or as the expectation over the estimated initial state distribution):

$$\hat{v}_{\text{AM}}(\pi) = \frac{1}{m} \sum_{j=1}^m \hat{V}_{\text{AM}}^{\pi}(s_1^{(j)}) \text{ or } \sum_{s_1 \in \mathcal{S}} \hat{\mu}_0(s_1) \hat{V}_{\text{AM}}^{\pi}(s_1),$$

$$\text{where } \hat{V}_{\text{AM}}^{\pi}(s) = \mathbb{E}_{\substack{a_t \sim \pi(\cdot | s_t) \\ r_t \sim \hat{r}(s_t, a_t) \\ s_{t+1} \sim \hat{p}(\cdot | s_t, a_t)}} \left[\sum_{t=1}^H \gamma^{t-1} r_t \mid s_1 = s \right].$$

Hyperparameters. When Monte-Carlo rollouts are used for evaluation, in lieu of a termination model (which itself requires specifying a threshold on the predicted termination probability), it is more common to directly set the maximum length of rollout H , which corresponds to the evaluation horizon. One possible heuristic is to set H based on the typical episode length \tilde{L} in historical data. However, if we expect the learned policy to outperform the behavior policy in a particular way and produce longer or shorter episodes (e.g., keeping patients healthy for a longer time, or discharge patients alive earlier), it may be more appropriate to adjust H depending on the specific setting. When analytic or iterative policy evaluation is used (for tabular settings), we consider $H \rightarrow \infty$.

Dependencies on Auxiliary Models. As mentioned above, AM depends on the following auxiliary models: \hat{p} , \hat{r} , and possibly $\hat{\mu}_0$, \hat{p}_{term} .

- For discrete state spaces, these quantities can be found using maximum likelihood estimation: $\hat{p}(s'|s, a) = \frac{\text{count}(s, a, s')}{\text{count}(s, a)}$, $\hat{r}(s, a) = \frac{\sum_{i=1}^N r_i \mathbb{1}(s_i=s, a_i=a)}{\text{count}(s, a)}$, $\hat{\mu}_0(s) = \frac{\text{count}(s_1=s)}{m}$ (the termination model can be considered part of the transition model).
- For continuous state spaces, these models can be found using supervised learning. Even though the true environment may be stochastic and one could model the full probability distributions of transitions and rewards in low-dimensional settings (e.g., Gaussian dynamics network (Torabi et al., 2018; Yu et al., 2020; Kidambi et al., 2020)), it makes the Monte-Carlo rollouts more susceptible to noise. One could generate multiple rollouts from each state, but this requires a new hyperparameter and significantly more computation. Alternatively, it is common to impose simplifying assumptions and instead learn the following regression models (Voloshin et al., 2019; Rajeswaran et al., 2020): $\hat{\Delta}_s(\mathbf{x}(s), a)$ an autoregressive model to predict the expected change in state features and $\hat{r}(\mathbf{x}(s), a)$ to predict the expected immediate reward. These models can be trained with mean squared error as the loss function. Furthermore, instead of modeling the full probability distribution of initial states as $\hat{\mu}_0$, we can directly use the observed initial states as the starting states for rollouts. Termination models are not typical, and instead we use H as the fixed rollout length.

Computational Cost. Similar to WIS, the cost of auxiliary model fitting does not depend on K , the number of candidate policies; only a single set of models needs to be learned using the following pattern sets: $\{\langle \mathbf{x}(s_i), a_i \rangle, \mathbf{x}(s'_i) - \mathbf{x}(s_i)\}_{i=1}^N$ for $\hat{\Delta}_s(s, a)$, and $\{\langle \mathbf{x}(s_i), a_i \rangle, r_i\}_{i=1}^N$ for $\hat{r}(s, a)$. The fitting complexity is thus $O(N)$. However, unlike WIS, the inference cost for AM scales linearly with K , because each candidate policy requires a different set of rollouts. The overall inference cost is $O(HmK)$, for advancing H steps from the m initial states following each of the K candidate policies, assuming they are all deterministic. The inference cost for a set of probabilistic policies is approximately $O(|\mathcal{A}|^H HmK)$ in the worst case, because we need to consider all possible action sequences with nonzero probability following each initial state and weight them appropriately.

3.2.3. FITTED Q EVALUATION (FQE)

Instead of reweighting the observed episodes (WIS) or building an approximate environment model (AM), we can algorithmically summarize the performance of the evaluation policy using the Q-function, a mapping from state-action pairs to the expected cumulative reward. This can be done by running FQE, an off-policy value-based algorithm that directly learns $\hat{Q}_{\text{FQE}}^\pi(s, a)$, the Q-function of π , from historical data.

Formalization. FQE (Alg. A.2) is a value-based, temporal difference algorithm (Le et al., 2019) and is closely related to expected SARSA with respect to a fixed policy for policy evaluation (van Seijen et al., 2009). Each iteration of the algorithm applies the Bellman equation for Q^π to compute the bootstrapping targets for all transitions from \mathcal{D} using the previous estimates (line 4) and then applies function approximation via supervised learning (lines 5-6). The algorithm outputs $\hat{Q}_{\text{FQE}}^\pi(s, a)$, the estimated Q-function of π . The final value estimate can be calculated as the average value for the initial states in \mathcal{D} :

$$\hat{v}_{\text{FQE}}(\pi) = \frac{1}{m} \sum_{j=1}^m \hat{V}_{\text{FQE}}^\pi(s_1^{(j)}), \text{ where } \hat{V}_{\text{FQE}}^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) \hat{Q}_{\text{FQE}}^\pi(s, a).$$

Hyperparameters. Similar to AM, FQE also has a hyperparameter H , the number of FQE iterations, which corresponds to the evaluation horizon (the output \tilde{Q}_H is the H -step Q-function of π).

Dependencies on Auxiliary Models. For each candidate policy π_k , we need to fit a sequence of Q-networks $\{\tilde{Q}_h\}_{h=1}^H$ (Alg. A.2, line 6).

Computational Cost. Unlike WIS and AM where the number of auxiliary models does not depend on the number of candidate policies (the fitting costs are amortized), FQE scales linearly with K because a different set of models is needed for each π_k . Furthermore, due to bootstrapping where the target value of the current iteration depends on the result of the previous iteration (Alg. A.2, line 4), each FQE iteration involves both inference and fitting of \tilde{Q} , resulting in a total fitting complexity of $O(HNK)$ for both deterministic and probabilistic candidate policies. Notably, the process of applying FQE to each policy cannot be parallelized due to this dependency. However, multiple policies can be evaluated in parallel. To obtain the final value estimate, each \tilde{Q}_H network is applied to the initial states from m episodes, resulting in inference complexity of $O(mK)$. Overall, FQE is much more expensive compared to WIS/AM in terms of computational cost (and more expensive than policy learning) and can take much longer to run.

3.2.4. WEIGHTED DOUBLY ROBUST (WDR) ESTIMATOR

WDR is a hybrid method that combines importance sampling from WIS with value estimators \widehat{V}^π and \widehat{Q}^π obtained from AM or FQE (Jiang and Li, 2016; Thomas and Brunskill, 2016). In addition to reweighting the observed returns (as in WIS), WDR leverages the value estimates by incorporating them at every time step with the hope of reducing the overall variance.

Formalization. We consider two versions of this estimator: WDR-AM and WDR-FQE. The estimator is defined as follows (with $\rho_{1:0} = 1$):

$$\hat{v}_{\text{WDR}}(\pi) = \frac{1}{m} \sum_{j=1}^m G_{\text{WDR}}(\tau_j; \pi),$$

$$\text{where } G_{\text{WDR}}(\tau_j; \pi) = \sum_{t=1}^{L^{(j)}} \left[\frac{\rho_{1:t}^{(j)}}{w_t} \gamma^{t-1} r_t^{(j)} - \left(\frac{\rho_{1:t}^{(j)}}{w_t} \widehat{Q}^\pi(s_t^{(j)}, a_t^{(j)}) - \frac{\rho_{1:t-1}^{(j)}}{w_{t-1}} \widehat{V}^\pi(s_t^{(j)}) \right) \right].$$

WDR is often assumed to be better than WIS and FQE/AM used by themselves because in theory, it is a consistent low-variance estimator if either importance ratios (i.e., behavior policy) or the value function estimates are properly specified (hence the name “doubly robust”). However, more recent empirical work has shown that WDR methods can be less accurate and have larger errors than its constituent OPE methods due to various reasons, such as stochastic transitions/rewards or high variance in importance ratios (Jiang and Li, 2016; Thomas and Brunskill, 2016; Voloshin et al., 2019).

Hyperparameters. WDR has the hyperparameters of its constituent parts (both ε and H).

Dependencies on Auxiliary Models. Beyond the models from WIS/AM/FQE, the two WDR approaches do not rely on any additional auxiliary models.

Computational Cost. There is no additional computation needed for model fitting. However, there are some notable differences in terms of inference complexity because WDR requires estimates of both Q and V for all N state-action pairs, rather than the m initial states.

- AM estimates Q/V by performing Monte-Carlo rollouts. Given the softened policy $\tilde{\pi}$, in order to obtain $\widehat{V}_{\text{AM}}^{\tilde{\pi}}(s) = \sum_{a \in \mathcal{A}} \tilde{\pi}(a|s) \widehat{Q}_{\text{AM}}^{\tilde{\pi}}(s, a)$ for each state, we need to perform rollouts that consider all possibilities of the future following $\tilde{\pi}$, which is exponential in the rollout length H . This results in $O(|\mathcal{A}|^H HNK)$ inference complexity to compute WDR-AM exactly, since the exponential rollouts need to be done for all N transitions and all K candidate policies, which is prohibitively expensive. If the evaluation policy is deterministic, then a slightly more practical solution is to use the softened $\tilde{\pi}$ to calculate importance ratios, but use the original π to perform rollouts, resulting in $O(HNK)$ inference complexity, which is still relatively expensive. In light of these computation requirements, we omitted WDR-AM from our experiments on problems with continuous state spaces in favor of WDR-FQE.
- Since FQE already learns parameterized functions for \widehat{Q} (and thus for \widehat{V}), we only need to perform additional inference runs using all N state-action pairs as input, resulting in $O(NK)$ inference complexity.

Table 1: Qualitative comparison of OPE methods when used for model selection over K candidate policies, including: any additional hyperparameters that must be set *a priori*, dependencies on additional estimated quantities (discrete state space) or auxiliary supervised models (continuous state space), and approximate computational complexity. For hybrid approaches (WDR), the listed complexity estimates are for additional computations not already done by the constituent OPE estimators. We assume the dataset consists of m episodes and N transitions, and all candidate policies are deterministic. ε is the policy softening hyperparameter and H is the length of the evaluation horizon. Other notation are described in Section 3.2. A more detailed breakdown can be found in Appendix B.

OPE	Hyperparameters	Dependencies	Complexity	
			Fitting	Inference
WIS	ε	$\hat{\pi}_b$	$O(N)$	$O(N)$
AM	H	\hat{p}, \hat{r}	$O(N)$	$O(HmK)$
FQE	H	$\{\tilde{Q}_h\}_{h=1}^H$ for each π_k	$O(HNK)$	$O(mK)$
WDR-AM	$\varepsilon; H$	WIS + AM	WIS + AM	$O(\mathcal{A} ^H HNK)$
WDR-FQE	$\varepsilon; H$	WIS + FQE	WIS + FQE	$O(NK)$

3.3. Combining Multiple OPE Estimators for Model Selection

While every OPE outputs an estimated value of the evaluation policy, each method does so in a different way, e.g., by estimating the behavior policy, estimating the environment model, or directly estimating the Q function. Based on past OPE literature (that focused on accuracy in estimating policy value (Voloshin et al., 2019; Fu et al., 2021)) and our analyses on computational costs (Table 1), we note that WIS/AM are more efficient but tend to be less accurate, whereas FQE is less efficient but tends to be more accurate. Given a large candidate set, it might not be worthwhile applying expensive OPE methods on low-quality policies. With this insight, we propose a simple, practical procedure for combining two OPE estimators.

Two-stage selection. Instead of selecting the best policy based on ranking from a single OPE method, we do this in two stages: in the first stage, a computationally efficient OPE (e.g., WIS) is applied to the entire candidate set in order to identify a promising subset (e.g., the top percentiles); in the second stage, a more accurate OPE (e.g., FQE) is used to rank within the pruned subset. The performance of the two-stage selection procedure depends on the initial subset size α as well as choices of the two OPE estimators. This additional hyperparameter α needs to be prespecified; in practice, this can be determined based on available computational resources. α dictates the total computational cost as well as to what extent the final selected policy depends on each of the two OPE estimators: $\alpha = 1$ defaults to the first OPE estimator, while $\alpha = K$ defaults to the second OPE estimator. See Appendix F for a detail discussion and theoretical analyses.

As baselines for comparison, we consider **average score** and **average ranking**, two naive ways of combining OPE scores that could take advantage of multiple OPE methods, but do not benefit from the computational savings as the two-stage approach and may be susceptible to outliers due to high variance of certain OPE estimators. We also compare to WDR (Section 3.2.4), a hybrid method that combines WIS with AM/FQE.

4. Experimental Setup

To test the utility of different OPE methods in model selection for offline RL, we performed a series of empirical experiments and evaluated how well OPE estimates of validation performance can help in selecting policies among a candidate set. All of our experiments rely on the following steps: (i) generating historical data $\mathcal{D}_{\text{train}}$ and \mathcal{D}_{val} from a simulator, (ii) applying an RL algorithm on $\mathcal{D}_{\text{train}}$ and obtaining the set of candidate policies by using different hyperparameters during training, (iii) using OPE and \mathcal{D}_{val} to estimate validation performance for each candidate policy, and (iv) quantifying the ranking quality and regret of OPE estimators with respect to true values of candidate policies. We discuss the details of these steps below.

4.1. Environment & Dataset Generation

While this work focuses on the offline setting, we consider a simulated problem because we need to evaluate the quality of policy ranking according to each OPE with respect to the true policy performance. We use the sepsis simulator adapted from Oberst and Sontag (2019) and Futoma et al. (2020a), which is modeled after the physiology of patients with sepsis. More details of the simulation setup can be found in Appendix C and the source code; below we provide a brief description. The underlying patient state consists of 5 discrete variables: a binary indicator for diabetes status, and 4 ordinal-valued vital signs (heart rate, blood pressure, oxygen concentration, glucose). We consider two state representations: a discrete state space with $|\mathcal{S}| = 1,440$, and a continuous state space where each state corresponds to a feature vector $\mathbf{x}(s) \in \{0, 1\}^{21}$ that uses a one-hot encoding for each underlying variable. There are 8 actions based on combinations of 3 binary treatments: antibiotics, vasopressors, and mechanical ventilation, where each treatment can affect certain vital signs and may raise/lower their values with pre-specified probabilities. A patient is discharged only when all vitals are normal and all treatments have been stopped; death occurs if 3 or more vitals are abnormal. Rewards are sparse and only assigned at the end of each episode, with +1 for survival and -1 for death, after which the system transitions into the respective absorbing state. For data generation, episodes are truncated at a maximum length of 20. We used a uniformly random policy to generate 10 pairs of datasets (for training and validation), each with large number of episodes ($m = 10,000$) and a different random seed. While this data generation process ensures sufficient exploration in the state-action space, it makes certain assumptions that might not hold in practice; we explore some of these issues in Section 5.2.

4.2. Implementation Details

To learn the policies, we ran FQI (Alg. A.1) for up to 50 iterations using a neural network function approximator. We used 10,000 episodes following the uniformly random policy for both training and validation unless otherwise specified, and repeated each experiment for 10 replication runs. We followed the procedures outlined in Section 3.2 to learn appropriate auxiliary models (either via MLE or as individual supervised tasks) and applied OPE methods using the learned models. We used $\varepsilon = 0.01$ and $H = 20$ as the default OPE hyperparameters following recommendations provided in Section 3.2. By default, we used neural networks consisting of one hidden layer with 1,000 neurons and ReLU activation to allow

for function approximators with sufficient expressivity. We trained these networks using the Adam optimizer (default settings) (Kingma and Ba, 2015) with a batch size of 64 for a maximum of 100 epochs, applying early stopping on 10% “validation data” (specific to each supervised task) with a patience of 10 epochs. We minimized either the mean squared error (MSE) for regression tasks (each iteration of FQI/FQE, dynamics/reward models of AM), or the cross-entropy loss with softmax activation for classification tasks (behavior policy for WIS). For FQI/FQE, we also added value clipping (to be within the range of possible returns $[-1, 1]$) when computing bootstrapping targets to ensure a bounded function class and encourage better convergence behavior (Mnih et al., 2015).

4.3. Evaluation

Given our goal is to select the best (or a good) policy from the candidate set, the OPE validation scores does not necessarily need to match exactly with true policy performance; they only need to correlate well (Irpan et al., 2019). Thus, we do not present mean squared error (MSE), a common metric used to evaluate OPE methods (Voloshin et al., 2019). Instead, we report the following quantitative metrics (Paine et al., 2020; Fu et al., 2021):

- Spearman’s rank correlation ρ between OPE scores and ground-truth policy values. This measures the overall ranking quality, i.e., how well an OPE method can rank policies according to their true performance from high to low.
- Regret, defined as $\max_{k=1 \dots K} v(\pi_k) - v(\pi_{k^*})$ where $k^* = \arg \max_{k'=1 \dots K} \hat{v}(\pi_{k'})$ corresponds to the candidate policy with the best OPE validation score. This measures how far the identified “best” policy is from the actual best policy among the candidate set. We also report top- n regret (i.e., $\text{regret}@n$), which is the smallest regret for the n policies with the highest OPE scores.

We evaluated policies analytically using the matrix inversion method and ground-truth MDP model parameters p , r , and μ_0 from the sepsis simulator to obtain the true value $v(\pi)$. While regret ultimately determines how well model selection can identify the best policy, we consider rank correlation and $\text{regret}@n$ as they indicate to what extent an OPE method can be used for identifying the promising subset in the two-stage selection.

5. Experiments & Results

In this section, we present empirical experiments that compare the utility of OPE methods for model selection in offline RL. In Section 5.1, we first consider the effectiveness of OPE methods in selecting neural network architectures and training hyperparameters. Specifically, we quantitatively measure the overall ranking quality, as well as the regret of selected policies. We also compare strategies for combining multiple OPE estimators, including naive averaging (of scores or ranking), WDR, as well as the proposed two-stage selection. To test the robustness of each OPE method in different settings and to enhance the generalizability of our findings, in Section 5.2, we explore robustness to the additional OPE hyperparameters as well as various data conditions with different sample sizes or different behavior policies. The main experiments focus on the continuous state space setting; results for the discrete state space setting can be found in Appendix D.2.

5.1. Model Selection of Neural Architectures & Training Hyperparameters

When a function approximator (such as neural networks) is used, there are many hyperparameters related to architectures and training procedures that affect model complexity and can lead to overfitting. In this experiment, we consider learning policies on the sepsis simulator with a wide range of hyperparameters (Table 2). Based on the hyperparameter combinations, we trained a total of 96 Q-networks using FQI on the training set, corresponding to 96 candidate policies, and then evaluated each candidate policy using OPE on validation data.

Table 2: Hyperparameter values considered for neural FQI on the sepsis simulator datasets.

Hyperparameter	Values
Hidden layers	{1, 2}
Hidden units	{100, 200, 500, 1000}
Learning rate	{1e-3, 1e-4}
FQI training iterations	{1, 2, 4, 8, 16, 32}

5.1.1. COMPARISON ACROSS OPE METHODS

To measure the ranking quality of each OPE method, we first inspect the scatter plots of OPE scores and true policy values (Figure 2 top). We quantify ranking quality using Spearman’s rank correlation ρ (reported as the means and standard deviations over 10 runs). In addition to overall ranking quality, we also measure regret@1, regret@5, and regret@10 to quantify how well model selection can help us choose the best (or a good) policy among the candidate set (Figure 2 bottom).

Results & Discussion. We observe that FQE produces the best ranking, with the highest $\rho = 0.89$ and points lying closest to the 45-degree line. WIS scores have a large variance and span almost the entire $[-1, 1]$ range even for policies with a true value > 0 , but there appears to be good separation between policies with value > 0 and those < 0 , leading to a relatively high $\rho = 0.74$. AM consistently underestimates the policy value, assigning scores close to 0 for almost all policies, yet it maintains a reasonable overall ranking with $\rho = 0.66$. The hybrid method WDR-FQE, though commonly assumed to outperform its constituent OPE estimators (in terms of MSE), actually produces worse ranking than both WIS and FQE, resulting in a lower $\rho = 0.64$. The poor accuracy (contributing to poor ranking) of hybrid methods has also been observed for other RL tasks in past work (Voloshin et al., 2019; Fu et al., 2021) and is likely due to various factors including limited data, environment stochasticity, estimated behavior policy, and model misspecification. We observe similar trends in regret@1 where FQE and WIS consistently achieves the lowest regret, while WDR-FQE sometimes leads to a very high regret (> 0.5). However, all methods achieve near-zero regret for regret@5 and regret@10 and in general can identify the best candidate policy in the top-5/top-10 set. We also note that validation scores based on FQI value estimates or the temporal difference error produce poor ranking with $\rho < 0.3$ and high regret (Appendix D.1), highlighting the need to use OPE for estimating validation performance.

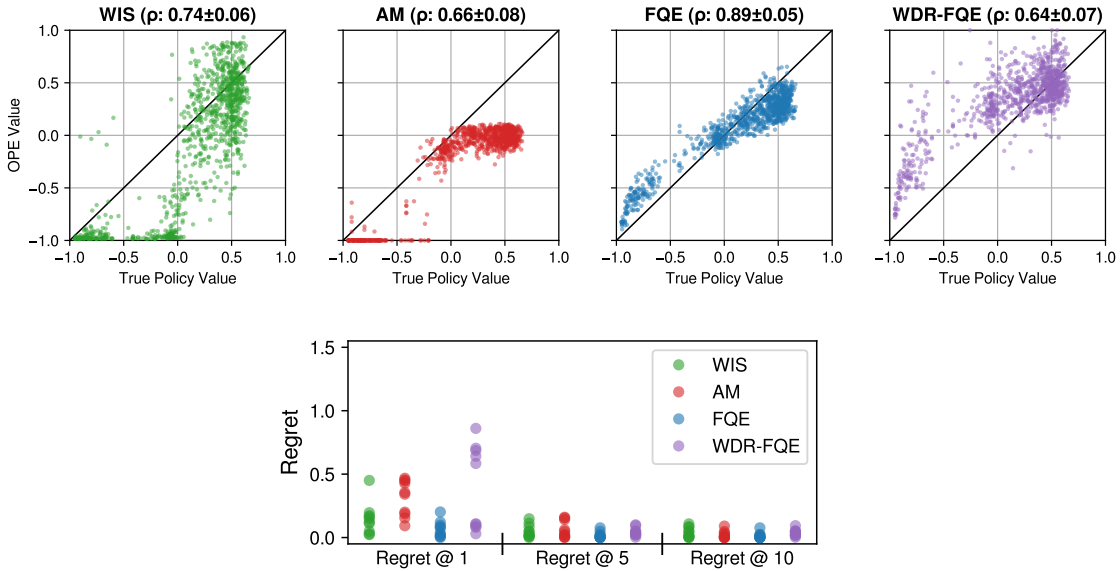


Figure 2: Top – Scatter plots comparing true policy value with OPE validation score. The results from 10 runs (each with 96 points) are overlaid in each OPE subplot. All values are clipped to the range of possible returns $[-1, 1]$ for clarity of visualization. Spearman’s rank correlations are listed in the title of each subplot (mean \pm std over 10 runs). FQE achieves the highest ρ and produces the best ranking over candidate policies. Bottom – Top- n regret of different OPE methods in model selection compared to the best candidate policy. While the best candidate policy does not always have the highest OPE validation score, it is frequently among those with the top-5/top-10 validation performance.

5.1.2. COMBINING MULTIPLE OPE ESTIMATORS

Given that WIS and FQE performed well, we explored various strategies of combining the two methods, in addition to WDR-FQE. As described in Section 3.3, we apply the proposed two-stage approach where WIS is used to select a promising subset (containing policies with the top WIS scores). In practice, the subset size α needs to be pre-determined (for example, based on computational constraints), but here we used a range of different sizes and explored their effect. We also consider two naive approaches – average score and average ranking of WIS and FQE – as baselines. We computed regret and repeated the experiments for the same 10 replication runs as before (Figure 3 left).

Results & Discussion. Simple strategies such as average score and average ranking outperform the more sophisticated WDR-FQE, achieving lower regret compared to WIS (and comparable regret with FQE). However, these approaches rely on computing both OPE estimators on the entire candidate set and are thus computationally expensive. The two-stage approach achieves comparable or lower regret than FQE across different subset sizes, giving the lowest regret at a subset size of 24 (and 48). Moreover, with minimal (but arguably necessary) additional computation compared to using WIS alone, we can achieve lower regret using the two-stage approach compared to using FQE alone, with potentially greater savings in computational costs (Figure 3 right).

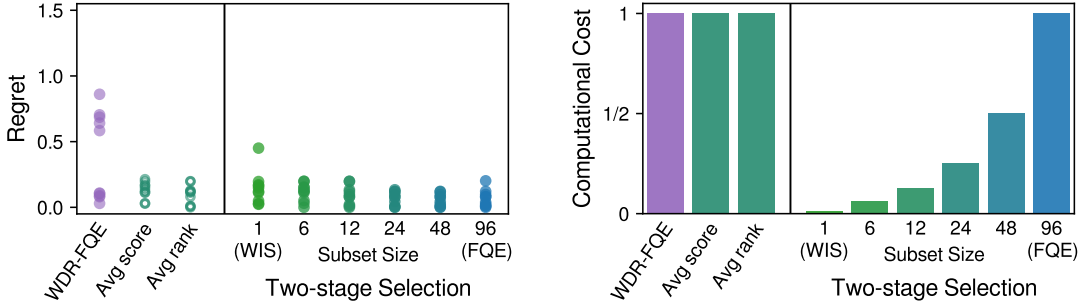


Figure 3: Comparison of regret (left) and relative computational cost (right) of different strategies of combining WIS and FQE for model selection. The computational costs are normalized with respect to a complete FQE evaluation over the entire candidate set ($K = 96$) and are for illustration purposes only. Two-stage selection with an intermediate subset size (e.g., 24) provides the best balance between performance and computational costs.

5.2. Sensitivity Analyses

In the experiments above, we made several assumptions about the validation data: (i) we used default auxiliary hyperparameters, (ii) we assumed a uniformly random behavior policy, and (iii) our validation data had a large sample size. These assumptions might not hold in practice. In this section, we vary the experimental conditions along these dimensions to test the generalizability of our conclusions. For these experiments, we kept the training set and training procedures the same throughout this comparison to generate the same candidate set for each run, so that the only difference comes from different validation data. We also kept the same neural network architecture and training procedures for auxiliary models. These experiments are repeated for the same 10 replication runs. We focus our results on the three non-hybrid OPE methods (WIS, AM, and FQE) as well as the proposed two-stage approach combining WIS with FQE with a subset size of 24. Additional sensitivity analyses for both discrete and continuous state spaces can be found in [Appendix D.2](#).

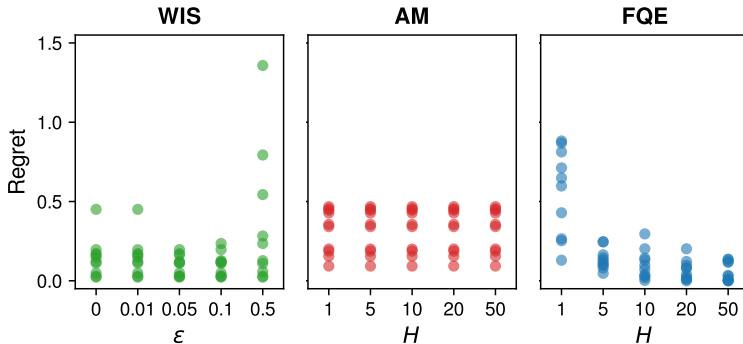


Figure 4: Regret comparisons for model selection with different auxiliary OPE hyperparameters. AM is the most robust to H but underperforms the other two OPE methods. WIS and FQE produce low regret given reasonable ϵ or H values.

5.2.1. SENSITIVITY TO OPE HYPERPARAMETERS

Since each OPE method has its own auxiliary hyperparameters, it is important to understand how they may affect OPE validation ranking quality, especially as there is no effective way of selecting them. In this experiment, we focus on the hyperparameters that must be set *a priori*, namely: the policy softening parameter ϵ for WIS and the evaluation horizon H for AM/FQE. We vary these hyperparameters in the reasonable ranges and measure the corresponding regret following model selection. For continuous state space models, even though the auxiliary models contain hyperparameters themselves (e.g., the neural network architecture), these can be tuned via standard model selection within each auxiliary supervised learning task; for consistency, we kept these hyperparameters the same throughout.

Results and Discussion. We visualize the regret in Figure 4. Among the three OPE methods, AM appears to be the most robust to different H values, although the regret it obtains is higher than the other two OPE methods and does not decrease as H increases. For WIS, the lowest regret is achieved at $\epsilon = 0.1$, but using ϵ values that are too small (potential issues with effective sample size) or too large ($\tilde{\pi}$ is too far from π) led to regret with slightly larger variance. FQE achieves more stable regret at larger evaluation horizons, e.g., $H \geq 10$, that correspond to a closer approximation to a full, infinite-horizon evaluation. Although WIS and FQE are more sensitive to their auxiliary hyperparameters, the performance is optimal (or close to optimal) following existing rules of thumb and recommendations we provided in Section 3.2.

5.2.2. SAMPLE SIZE OF VALIDATION DATA

In our main experiments, we used validation data with 10,000 episodes, which is as large as the training data. In supervised learning, the size of validation data is more commonly only a fraction of the training data size. In this comparison, we reduced the amount of validation data used for OPE and measured the regret.

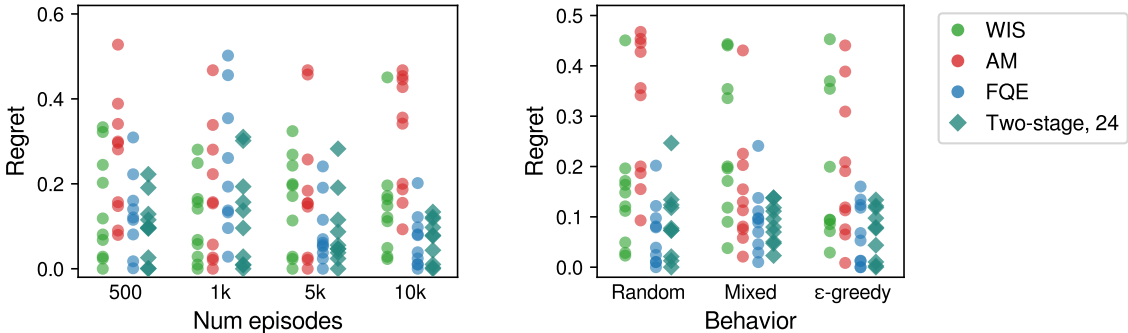


Figure 5: Regret comparison for OPE using validation data with varying sample sizes and varying behaviors. Left – A larger validation set generally results in lower regret across all OPE methods. Right – Less exploratory behavior can lead to higher regret. In almost all scenarios, the proposed two-stage approach is able to maintain low regret while reducing overall computational cost.

Results and Discussion. All OPE methods tend to benefit from larger validation sizes, achieving lower regret (Figure 5, left). The trend is less apparent for WIS (which generally results in noisier value estimates) and for AM (possibly due to model misspecification). The proposed two-stage selection (combining WIS and FQE with a subset size of 24) consistently matches (and sometimes slightly outperforms) FQE in terms of regret, except for very small sample sizes (e.g., 500 episodes). This could be because WIS is more susceptible to large variances due to limited sample sizes and produces less reliable pruning in the first stage. Given reasonably sized validation data (1,000 episodes or more for this problem), the proposed two-stage approach maintains performance comparable to FQE, while reducing computational costs.

5.2.3. TYPE OF BEHAVIOR & EXTENT OF EXPLORATION

The behavior policies used to collect validation data can affect which part of the state-action space receives sufficient exploration and the resultant OPE model selection performance. While we previously assumed a uniformly random policy that provides sufficient exploration, this is rarely the case in real life; e.g., clinicians do not behave randomly when selecting treatments. Moreover, historical data may be generated by multiple clinicians who use different policies. Here, we examine how the behavior policy affects model selection. We first considered validation data containing 10,000 episodes following an ϵ -greedy behavior policy (with $\epsilon = 0.1$); though this produces near-optimal behavior, it limits the extent of exploration. We also considered data that contained 5,000 episodes of both the random policy and the ϵ -greedy policy to create a mixture of behaviors (with the same total sample size of 10,000 episodes). Although these two settings violate the assumption that training data and validation data follow the same data distribution (as discussed in Section 3), the evaluation problem is still valid since all data are drawn from the same underlying MDP environment.

Results and Discussion. The random behavior policy produces more exploratory data compared to an ϵ -greedy policy, leading to slightly lower regret in general (Figure 5, right). Surprisingly, even when a mixture of multiple behavior policies is used, WIS is not heavily affected by this policy misspecification and still leads to low regret. Reassuringly, the proposed two-stage selection (with a subset size of 24) is able to maintain (or outperform) the regret achieved by FQE while reducing overall computation, even in challenging situations with less exploratory data or mixed behaviors.

6. Conclusion

Model selection for offline RL is challenging because we cannot easily obtain the equivalent of “validation performance” given only historical interaction data. To date, RL for healthcare research has skirted the issue of model selection by either assuming default hyperparameters or relying on ground-truth performance only obtainable with simulators. In this work, we explored a model selection framework for offline RL that uses OPE as a proxy estimate for validation performance. We provided an in-depth analysis of how this can be implemented using four popular OPE methods, detailing the practical considerations such as the reliance on auxiliary hyperparameters/models and computational requirements. On a simulated

RL task of sepsis treatment, we found that FQE consistently provides the most accurate ranking over candidate policies and is the most robust to various data conditions (sample sizes and behavior policies). Given the high computational cost of FQE, we proposed a simple two-stage approach to combine two OPE estimators and demonstrated that it outperforms simple averaging or more sophisticated doubly-robust methods, while avoiding expensive computation on low-quality policies.

There has been increasing interest on the problem of model selection for offline RL in recent years (Appendix E). Recently, Paine et al. (2020) demonstrated that FQE is a useful strategy for hyperparameter selection in offline RL on several benchmark control tasks. In contrast, we compare multiple OPE estimators both qualitatively and quantitatively to address questions of practical importance for healthcare settings. Our empirical results further corroborate the findings of Irpan et al. (2019) and Paine et al. (2020), who showed that FQI values and TD errors are poor validation metrics due to overestimation of value functions during training, thus highlighting the importance and necessity of using OPE to estimate validation performance. During the preparation of this work, Fu et al. (2021) published a benchmark suite for comparing OPE methods on continuous control tasks, with Spearman’s correlation as one of evaluation metrics. Similar to our empirical findings, they found FQE to consistently produce the best results, whereas other OPE approaches (including AM and WIS) varied more with respect to the specific task and modeling choice. However, their focus is on presenting a standardized benchmark to compare OPE methods, rather than carefully reflecting on the practical and implementation issues that arise when using OPE for model selection.

We note several important limitations of this work. First, our experiments are performed on a simulated RL task because of the need to obtain true performance of learned policies. Though we explored different data conditions including behavior policies and sample sizes, our simulation might not capture all nuances in real data collection. In particular, certain properties of the simulation (discrete action spaces, short horizons, noiseless observations, no missingness, sparse terminal rewards) may limit the generalizability of our conclusions to real observational data and other RL problems. Using a simulator also does not address important issues that may arise when designing the RL problem setup based on real datasets (e.g., defining a good action space or meaningful initial states). We encourage future work to investigate more healthcare RL problems using our released code. Though the simulated sepsis task may not directly produce any clinically relevant insights for treating sepsis, we expect future work on sepsis management (Komorowski et al., 2018; Raghu et al., 2017a,b; Killian et al., 2020) to build upon our lessons and improve their model selection procedure, thus enabling fairer comparisons. Second, we only considered four OPE methods that are commonly used in healthcare tasks. It could be beneficial to understand the pros and cons of other OPE methods when incorporated into the model selection framework, e.g., approaches based on marginalized importance sampling (Liu et al., 2018; Zhang et al., 2019; Uehara et al., 2020). Some of these techniques have been compared quantitatively in concurrent work by Fu et al. (2021) on a continuous control benchmark suite, but none of them outperformed FQE in terms of ranking quality. Third, we assumed the same validation data is used to fit the auxiliary models and calculate OPE estimates. It is possible to use two separate validation sets (one for fitting models, the other for applying OPE); we preferred our current strategy as it leverages more data for both steps and empirically produced

reasonable results, though this may violate the independence assumptions needed for the theoretical guarantees of OPE (Jiang and Li, 2016). We also did not implement additional “tricks” that could help improve OPE accuracy, e.g., truncated importance weights for WIS (Ionides, 2008; Swaminathan and Joachims, 2015) or more sophisticated models for AM (Torabi et al., 2018; Yu et al., 2020; Kidambi et al., 2020). Lastly, we generated datasets with sufficient exploration and coverage over the entire state-action space, and trained all policies using a vanilla RL algorithm (i.e., FQI) that does not address distribution shift. Recently proposed algorithms that specifically account for distribution shift (Kidambi et al., 2020; Yu et al., 2020; Kumar et al., 2020; Liu et al., 2020) require additional hyperparameters, and we expect our framework to be useful for selecting these hyperparameters if the OPE step also considers distribution shift (e.g., by restricting the state-conditional action space to disallow rare actions). Despite these limitations, our empirical results suggest that using OPE for model selection in offline RL is a promising solution.

In this paper, we outlined important practical considerations when using OPE for model selection in offline RL. Our work serves as a guide for RL practitioners in healthcare and other high-stakes domains when considering the model selection step, and can ultimately lead to better policies on real-world datasets. We encourage researchers in RL for health to build upon our lessons and report model selection techniques, in an effort to improve reproducibility, enable fair comparisons, and advance the field.

Acknowledgments

This work was supported by the National Science Foundation (NSF award no. IIS-1553146) and the National Library of Medicine (NLM grant no. R01LM013325). The views and conclusions in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the National Science Foundation nor the National Library of Medicine. This work was supported in part through computational resources and services provided by [Advanced Research Computing](#) at the University of Michigan, Ann Arbor. The authors would like to thank [Aditya Modi](#) and members of the [MLD3 group](#) for helpful discussions regarding this work, as well as the reviewers for their constructive feedback.

References

- James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(2), 2012.
- Xinyun Chen, Lu Wang, Yizhe Hang, Heng Ge, and Hongyuan Zha. Infinite-horizon off-policy policy evaluation with multiple behavior policies. In *International Conference on Learning Representations*, 2019.
- Christoph Dann, Gerhard Neumann, Jan Peters, et al. Policy evaluation with temporal differences: A survey and comparison. *Journal of Machine Learning Research*, 15:809–883, 2014.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556, 2005.

- Amir-massoud Farahmand and Csaba Szepesvári. Model selection in reinforcement learning. *Machine learning*, 85(3):299–332, 2011.
- M Fard and Joelle Pineau. PAC-Bayesian model selection for reinforcement learning. *Advances in Neural Information Processing Systems*, 23:1624–1632, 2010.
- Justin Fu, Mohammad Norouzi, Ofir Nachum, George Tucker, Alexander Novikov, Mengjiao Yang, Michael R Zhang, Yutian Chen, Aviral Kumar, Cosmin Paduraru, et al. Benchmarks for deep off-policy evaluation. In *International Conference on Learning Representations*, 2021.
- Scott Fujimoto, Edoardo Conti, Mohammad Ghavamzadeh, and Joelle Pineau. Benchmarking batch deep reinforcement learning algorithms. *arXiv preprint arXiv:1910.01708*, 2019a.
- Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pages 2052–2062. PMLR, 2019b.
- Joseph Futoma, Michael C Hughes, and Finale Doshi-Velez. POPCORN: Partially observed prediction constrained reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, 2020a.
- Joseph Futoma, Muhammad A Masood, and Finale Doshi-Velez. Identifying distinct, effective treatments for acute hypotension with SODA-RL: Safely optimized diverse accurate reinforcement learning. *AMIA Summits on Translational Science Proceedings*, 2020:181, 2020b.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- Omer Gottesman, Fredrik Johansson, Joshua Meier, Jack Dent, Donghun Lee, Srivatsan Srinivasan, Linying Zhang, Yi Ding, David Wihl, Xuefeng Peng, et al. Evaluating reinforcement learning algorithms in observational health settings. *arXiv preprint arXiv:1805.12298*, 2018.
- Josiah Hanna, Scott Niekum, and Peter Stone. Importance sampling policy evaluation with an estimated behavior policy. In *International Conference on Machine Learning*, pages 2605–2613. PMLR, 2019.
- Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- Edward L Ionides. Truncated importance sampling. *Journal of Computational and Graphical Statistics*, 17(2):295–311, 2008.
- Alex Irpan, Kanishka Rao, Konstantinos Bousmalis, Chris Harris, Julian Ibarz, and Sergey Levine. Off-policy evaluation via off-policy classification. In *Advances in Neural Information Processing Systems*, 2019.

- Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.
- Nan Jiang, Alex Kulesza, Satinder Singh, and Richard Lewis. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1181–1189. Citeseer, 2015.
- Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, pages 651–673, 2018.
- Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. MOReL: Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Taylor W Killian, Haoran Zhang, Jayakumar Subramanian, Mehdi Fatemi, and Marzyeh Ghassemi. An empirical study of representation learning for reinforcement learning in healthcare. In *Machine Learning for Health*, pages 139–160. PMLR, 2020.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Matthieu Komorowski, Leo A Celi, Omar Badawi, Anthony C Gordon, and A Aldo Faisal. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11):1716, 2018.
- Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy Q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, pages 11784–11794, 2019.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative Q-learning for offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Ilja Kuzborskij, Claire Vernade, András György, and Csaba Szepesvári. Confident off-policy evaluation and selection through self-normalized importance weighting. In *International Conference on Artificial Intelligence and Statistics*, pages 640–648. PMLR, 2021.
- Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer, 2012.
- Romain Laroche, Paul Trichelair, and Remi Tachet Des Combes. Safe policy improvement with baseline bootstrapping. In *International Conference on Machine Learning*, pages 3652–3661. PMLR, 2019.
- Hoang Le, Cameron Voloshin, and Yisong Yue. Batch policy learning under constraints. In *International Conference on Machine Learning*, pages 3703–3712, 2019.

- Byungjun Lee, Jongmin Lee, Peter Vrancx, Dongho Kim, and Kee-Eung Kim. Batch reinforcement learning with hyperparameter gradients. In *International Conference on Machine Learning*, pages 5725–5735. PMLR, 2020.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, pages 661–670, 2010.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research*, 18(1):6765–6816, 2017.
- Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, 2018.
- Yao Liu, Adith Swaminathan, Alekh Agarwal, and Emma Brunskill. Provably good batch off-policy reinforcement learning without great exploration. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. Offline policy evaluation across representations with applications to educational games. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1077–1084, 2014.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- Michael Oberst and David Sontag. Counterfactual off-policy evaluation with Gumbel-max structural causal models. In *International Conference on Machine Learning*, pages 4881–4890, 2019.
- Cosmin Păduraru, Doina Precup, Joelle Pineau, and Gheorghe Comănici. An empirical analysis of off-policy learning in discrete mdps. In *European Workshop on Reinforcement Learning*, pages 89–102, 2013.
- Tom Le Paine, Cosmin Paduraru, Andrea Michi, Caglar Gulcehre, Konrad Zolna, Alexander Novikov, Ziyu Wang, and Nando de Freitas. Hyperparameter selection for offline reinforcement learning. *arXiv preprint arXiv:2007.09055*, 2020.

- Niranjani Prasad, Li-Fang Cheng, Corey Chivers, Michael Draugelis, and Barbara E Engelhardt. A reinforcement learning approach to weaning of mechanical ventilation in intensive care units. *33rd Conference on Uncertainty in Artificial Intelligence*, 2017.
- Aniruddh Raghu, Matthieu Komorowski, Imran Ahmed, Leo Celi, Peter Szolovits, and Marzyeh Ghassemi. Deep reinforcement learning for sepsis treatment. *NeurIPS 2017 ML4H Workshop*, 2017a.
- Aniruddh Raghu, Matthieu Komorowski, Leo Anthony Celi, Peter Szolovits, and Marzyeh Ghassemi. Continuous state-space models for optimal sepsis treatment: a deep reinforcement learning approach. In *Proceedings of the 2nd Machine Learning for Healthcare Conference*, volume 68, pages 147–163. PMLR, 2017b.
- Aravind Rajeswaran, Igor Mordatch, and Vikash Kumar. A game theoretic framework for model based reinforcement learning. In *International Conference on Machine Learning*, pages 7953–7963. PMLR, 2020.
- Martin Riedmiller. Neural fitted Q iteration – first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, pages 317–328. Springer, 2005.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- Jasper Snoek, Hugo Larochelle, and Ryan Prescott Adams. Practical bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 2012.
- Adith Swaminathan and Thorsten Joachims. The self-normalized estimator for counterfactual learning. *Advances in Neural Information Processing Systems*, 28, 2015.
- Shengpu Tang, Aditya Modi, Michael Sjoding, and Jenna Wiens. Clinician-in-the-loop decision making: Reinforcement learning with near-optimal set-valued policies. In *International Conference on Machine Learning*, pages 9387–9396. PMLR, 2020.
- Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148. PMLR, 2016.
- Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4950–4957, 2018.
- Masatoshi Uehara, Jiawei Huang, and Nan Jiang. Minimax weight and Q-function learning for off-policy evaluation. In *International Conference on Machine Learning*, pages 9659–9668. PMLR, 2020.

- Harm van Seijen, Hado van Hasselt, Shimon Whiteson, and Marco Wiering. A theoretical and empirical analysis of Expected Sarsa. In *2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, pages 177–184, 2009.
- Cameron Voloshin, Hoang M Le, Nan Jiang, and Yisong Yue. Empirical study of off-policy policy evaluation for reinforcement learning. *arXiv preprint arXiv:1911.06854*, 2019.
- Tengyang Xie and Nan Jiang. Batch value-function approximation with only realizability. In *International Conference on Machine Learning*, 2021.
- Chao Yu, Jiming Liu, and Shamim Nemati. Reinforcement learning in healthcare: A survey. *arXiv preprint arXiv:1908.08796*, 2019.
- Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. MOPO: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Ruiyi Zhang, Bo Dai, Lihong Li, and Dale Schuurmans. GenDICE: Generalized offline estimation of stationary values. In *International Conference on Learning Representations*, 2019.

Appendix A. Pseudocode for FQI & FQE

Below we present the pseudocode for FQI (Alg. A.1) and FQE (Alg. A.2), using notation we introduced in Section 2.

Algorithm A.1 Fitted Q Iteration
(Ernst et al., 2005; Riedmiller, 2005)

```

1: Input: dataset  $\mathcal{D}$ ,
   function class  $\mathcal{F}$ , max iteration  $H$ 
2: Initialize  $\tilde{Q}_0(s_i, a_i) = 0$  for all  $(s_i, a_i) \in \mathcal{D}$ 
3: for  $h$  in  $1 \dots H$  do
4:   Compute target  $y_i = r_i + \gamma \max_{a' \in \mathcal{A}} \tilde{Q}_{h-1}(s'_i, a') \forall i$ 
5:   Build training set  $\tilde{\mathcal{D}}_h = \{(s_i, a_i), y_i\}_{i=1}^N$ 
6:   Solve a supervised learning problem
        $\tilde{Q}_h \leftarrow \arg \min_{f \in \mathcal{F}} \sum_{i=1}^N (f(s_i, a_i) - y_i)^2$ 
7: end for
8: Output:  $\pi(s) = \arg \max_{a \in \mathcal{A}} \tilde{Q}_H(s, a)$ 

```

Algorithm A.2 Fitted Q Evaluation
(Le et al., 2019; Voloshin et al., 2019)

```

1: Input: dataset  $\mathcal{D}$ , policy to be evaluated  $\pi$ ,
   function class  $\mathcal{F}$ , max iteration  $H$ 
2: Initialize  $\tilde{Q}_0(s_i, a_i) = 0$  for all  $(s_i, a_i) \in \mathcal{D}$ 
3: for  $h$  in  $1 \dots H$  do
4:   Compute target  $y_i = r_i + \gamma \sum_{a' \in \mathcal{A}} \pi(a'|s'_i) \tilde{Q}_{h-1}(s'_i, a') \forall i$ 
5:   Build training set  $\tilde{\mathcal{D}}_h = \{(s_i, a_i), y_i\}_{i=1}^N$ 
6:   Solve a supervised learning problem
        $\tilde{Q}_h \leftarrow \arg \min_{f \in \mathcal{F}} \sum_{i=1}^N (f(s_i, a_i) - y_i)^2$ 
7: end for
8: Output:  $\tilde{Q}_H$ 

```

Appendix B. OPE Computational Requirements

As we have mentioned in Section 3.2, each OPE estimator depends on auxiliary models. For example, WIS requires a model of the behavior policy $\hat{\pi}_b$, whereas AM requires the transition/reward models \hat{p} and \hat{r} . Here, we further discuss how these models can be obtained from data, for RL problems with either a discrete or continuous state space. We assume a discrete action space and leave continuous action spaces to future work.

Discrete state space. This corresponds to the tabular setting, where the “models” are probability/value tables and can be obtained via maximum likelihood estimation (MLE) using count-based averaging:

$$\begin{aligned} \hat{\pi}_b(a|s) &= \frac{\text{count}(s, a)}{\text{count}(s)} & \hat{p}(s'|s, a) &= \frac{\text{count}(s, a, s')}{\text{count}(s, a)} \\ \hat{r}(s, a) &= \frac{\sum_{i=1}^N r_i \mathbb{1}_{(s_i=s, a_i=a)}}{\text{count}(s, a)} & \hat{\mu}_0(s) &= \frac{\text{count}(s_1 = s)}{m} \end{aligned}$$

where $\text{count}(s, a)$ and $\text{count}(s, a, s')$ are the (marginal) counts of the state-action pair and state-action-next-state transition in \mathcal{D} respectively. These quantities can be found relatively efficiently by looping through the dataset once and accumulating the respective counts.

Continuous state space. In this case, a state is represented as a feature vector $\mathbf{x}(s) \in \mathbb{R}^d$. Thus, all of the above-mentioned models need to be approximated using some parameterized functions (e.g., neural networks). These models (as shown in Table 3) are independent from the main policy learning process; they can be treated as classification or regression tasks and trained using standard supervised learning techniques (e.g., stochastic gradient descent with early stopping on a validation split). Furthermore, each of these models may be used for inference for multiple times given different inputs. We summarize the total number of additional models and inference runs for each OPE in Table 4.

Table 3: Details of auxiliary supervised models required by each OPE for problems with continuous state spaces.

OPE	Model	Pattern Set	Task Type	Loss
		$\{x_i, y_i\}_{i=1}^N$		
WIS	$\hat{\pi}_b(a s)$	$\mathbf{x}(s_i)$	classification	cross entropy
AM	$\hat{\Delta}_s(s, a)$	$\langle \mathbf{x}(s_i), a_i \rangle$	regression	mean squared error
	$\hat{r}(s, a)$	$\langle \mathbf{x}(s_i), a_i \rangle$	regression	mean squared error
FQE	$\tilde{Q}_h(s, a)$	$\langle \mathbf{x}(s_i), a_i \rangle$	regression	mean squared error
		$r_i + \max_{a' \in \mathcal{A}} \tilde{Q}_{h-1}(s'_i, a')$		

Table 4: Summary of computational requirements of OPE used for model selection for K candidate policies. We assume the dataset consists of m episodes and N transitions, and H is the length of the evaluation horizon. For WDR we only list additional computation that are not already computed in other OPE estimators.

OPE	Fitting		Inference	
	Runs	Samples	Runs	Samples
WIS	1	N	1	N
AM	2	N	$2K$	mH
FQE	KH	N	$KH; K$	$N; m$
WDR-AM	—	—	$2 \mathcal{A} ^H HK$	N
WDR-FQE	—	—	K	N

Appendix C. Sepsis Simulator - Details

We use the sepsis simulator adapted from Oberst and Sontag (2019) and Futoma et al. (2020a), which is crudely modeled after the physiology of patients with sepsis. Our experiments are based on an implementation publicly available at <https://github.com/clinicalml/gumbel-max-scm/tree/sim-v2/sepsisSimDiabetes>.

- **Action space.** There are 8 actions based on combinations of 3 binary treatments: antibiotics, vasopressors, and mechanical ventilation, where each treatment can affect certain vital signs and may raise/lower their values with pre-specified probabilities.
- **State space.** The underlying patient state consists of 5 discrete variables: a binary indicator for diabetes status, and 4 ordinal-valued vital signs (heart rate, blood pressure, oxygen concentration, glucose). The previously administered treatments are also encoded in the state as 3 additional binary variables. See Table 5 for more details. We consider two state representations: a discrete state space with $|\mathcal{S}| = 1,440$, as well as a feature vector $\mathbf{x}(s) \in \{0, 1\}^{21}$ that uses a one-hot encoding for each underlying variable. We separately added two absorbing states to represent discharge and death.
- **Initial states distribution.** The diabetes indicator is set to 1 w.p. 0.2. To make the problem easier, the MDP can start from any non-terminating state (including those with treatments enabled) with equal probability (conditioned on the diabetes indicator value), for a total of $303 \times 2 = 606$ initial states.

- **Transition dynamics.** The treatments (either added or withdrawn) affect different state variables and are applied sequentially as detailed in Table 6.
- **Rewards & termination condition.** A patient is discharged only when all vitals are normal and all treatments have been stopped; death occurs if 3 or more vitals are abnormal. Rewards are sparse and only assigned at the end of each episode (when transitioning from a terminating state to one of the two absorbing states), with +1 for survival and -1 for death, after which the system transitions into the respective absorbing state and obtains 0 reward afterwards.

Table 5: Details on state variables.

Variable	Name	Levels (numeric / text)	
hr	Heart Rate	{0,1,2}	{L,N,H}
sbp	Systolic Blood Pressure	{0,1,2}	{L,N,H}
o2	Percent Oxygen	{0,1}	{L,N}
glu	Glucose	{0,1,2,3,4}	{LL,L,N,H,HH}
abx	Antibiotics	{0,1}	{off,on}
vaso	Vasopressors	{0,1}	{off,on}
vent	Mechanical Ventilation	{0,1}	{off,on}
diab	Diabetes Indicator	{0,1}	{no,yes}

Table 6: Details on transition dynamics.

Step	Variable	Current	New	Change	Effect
1	abx	-	on	-	hr H → N w.p. 0.5 sbp H → N w.p. 0.5
		on	off	withdrawn	hr N → H w.p. 0.1 sbp N → H w.p. 0.5
2	vent	-	on	-	o2 L → N w.p. 0.7
		on	off	withdrawn	o2 N → L w.p. 0.1
3	vaso	-	on	-	L → N w.p. 0.7 (non-diabetic) N → H w.p. 0.7 (non-diabetic) sbp L → N w.p. 0.5 (diabetic) L → H w.p. 0.4 (diabetic) N → H w.p. 0.9 (diabetic)
		on	off	withdrawn	glu LL → L, L → N, N → H, H → HH w.p. 0.5 (diabetic) sbp N → L w.p. 0.1 (non-diabetic) H → N w.p. 0.1 (non-diabetic) N → L w.p. 0.05 (diabetic) H → N w.p. 0.05 (diabetic)
4	hr	fluctuate			vitals spontaneously fluctuate when not affected by treatment (either enabled or withdrawn) • the level fluctuates ±1 w.p. 0.1, except: • glucose fluctuates ±1 w.p. 0.3 (diabetic)
5	sbp				
6	o2				
7	glu				

Appendix D. Additional Experimental Results

D.1. Model Selection using FQI Values & TD Errors

Following Section 5.1, we consider two additional metrics to be used as a proxies of validation performance.

- FQI values. This is based on the predicted values from the Q-network learned by FQI. It is in a way similar to “training performance” in the supervised learning setting. Similar to all OPE scores, we compute FQI score with respect to the empirical initial state distribution (from the validation set):

$$\hat{v}_{\text{FQI}}(\pi) = \frac{1}{m} \sum_{j=1}^m \hat{V}_{\text{FQI}}^{\pi}(s_1^{(j)}), \text{ where } \hat{V}_{\text{FQI}}^{\pi}(s) = \sum_{a \in \mathcal{A}} \pi(a|s) \hat{Q}_{\text{FQI}}^{\pi}(s, a).$$

- RMS-TDE, i.e., root-mean-squared temporal difference error (Dann et al., 2014), is the objective function (i.e., “loss function”) used in temporal difference learning algorithms such as FQI (more precisely, the loss function is MS-TDE without the square root). This is commonly calculated using the empirical distribution of transitions from the validation set.

$$\ell = \sqrt{\mathbb{E}_{(s_i, a_i, r_i, s'_i) \sim \mathcal{D}}[\delta_i^2]},$$

where $\delta_i = \left[r_i + \gamma \sum_{a'} \pi(a'|s'_i) \hat{Q}(s'_i, a') \right] - \hat{Q}(s_i, a_i)$ defines the temporal difference error of a transition tuple, and \hat{Q} is the same $\hat{Q}_{\text{FQI}}^{\pi}$ from FQI training.

We evaluate how well FQI values and RMS-TDE can serve as “validation scores” for model selection by measuring the ranking quality and regret (Figure 6). We observe these two “validation scores” clearly underperforms the OPE methods considered in the main paper, achieving Spearman’s rank correlation $|\rho| < 0.3$ with very little trend in the scatter plots, and high regret even when considering the top-5/top-10 set. This suggests that FQI values and RMS-TDE are not effective metrics for model selection in offline RL, corroborating the findings of Irpan et al. (2019) and Paine et al. (2020).

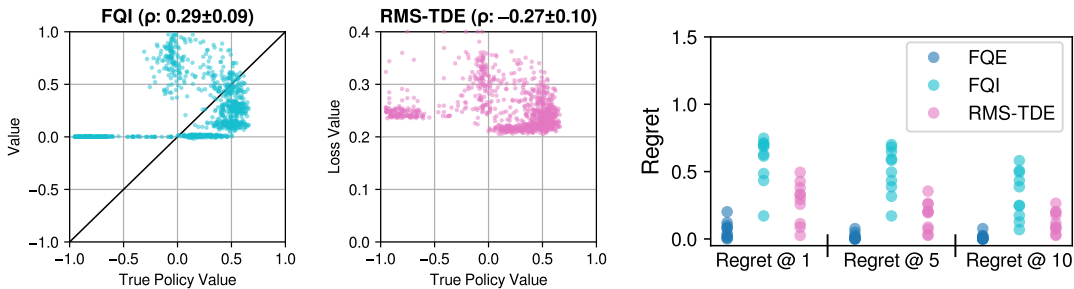


Figure 6: Left – Scatter plots comparing true policy value with FQI values and RMS-TDE. All values are clipped to the range of possible returns $[-1, 1]$ for clarity of visualization. Right – top- n regret in model selection for FQI and RMS-TDE.

D.2. Sensitivity Analyses for Early Stopping

In this section, we focus on the model selection problem of early stopping, i.e., identifying the best number of training iterations in FQI. This allows us to consider both tabular and function approximation settings. The candidate policy set is made up of all the policies derived from Q-functions along the training path of an FQI run. We focus our experiments on WIS/AM/FQE, because WDR resulted in large variances and heavily skews the plots.

D.2.1. SENSITIVITY TO OPE HYPERPARAMETERS

Since each OPE method has its own auxiliary hyperparameters, it is important to understand how they may affect OPE validation ranking quality, especially as there is no effective way of selecting them. In this experiment, we focus on the hyperparameters that must be set *a priori*, namely: the policy softening parameter ϵ and the evaluation horizon H for AM/FQE. We vary these hyperparameters in the reasonable ranges and measure the corresponding regret following model selection (similar to Section 5.2.1). For continuous state space models, even though the auxiliary models contain hyperparameters themselves (e.g., the neural architecture), these can be tuned via standard model selection within each auxiliary supervised learning task; for consistency, we kept these hyperparameters the same throughout. Note that in the tabular setting, AM can be computed analytically (instead of via Monte-Carlo rollouts) and does not have any hyperparameters, and is therefore omitted from this experiment.

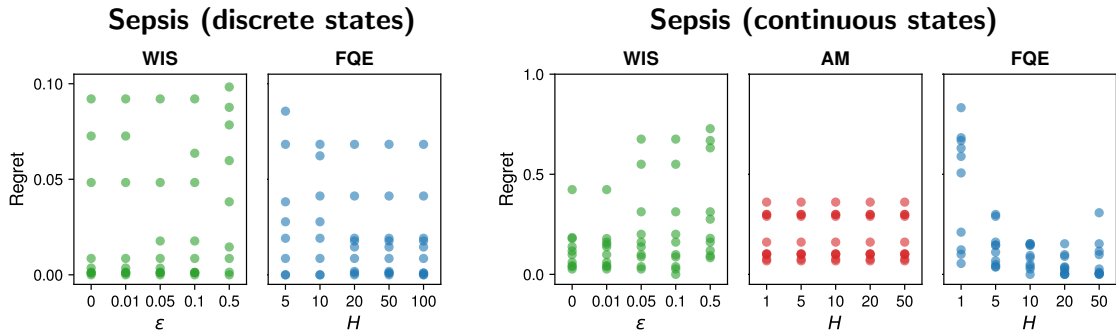


Figure 7: Regret comparisons for model selection with different OPE hyperparameters.

Results and Discussion. For the tabular setting (Figure 7 left), both WIS and FQE are relatively robust to hyperparameter values (ϵ and H , respectively) within the reasonable ranges, maintaining a low regret of < 0.1 . For the continuous state setting (Figure 7 right), although AM appears the most robust to different H values, the regret it obtains is higher than the other two OPE and does not decrease as H increases. WIS obtains the lowest regret at a low ϵ , where using ϵ values that are too large ($\tilde{\pi}$ is too far from π) led to regret with slightly larger variance. FQE achieves the lowest regret at an intermediate value of $H = 20$. This could be due to a form of overfitting within auxiliary models that could cause longer evaluation horizons to have poorer estimates when approximation errors are present (Jiang et al., 2015).

D.2.2. SIZE OF VALIDATION SET

In this experiment, we vary the amount of validation data used for OPE, similar to [Section 5.2.2](#). The training set and training procedures were kept the same to generate the same candidate set for each run, so the only difference is the validation sample size.

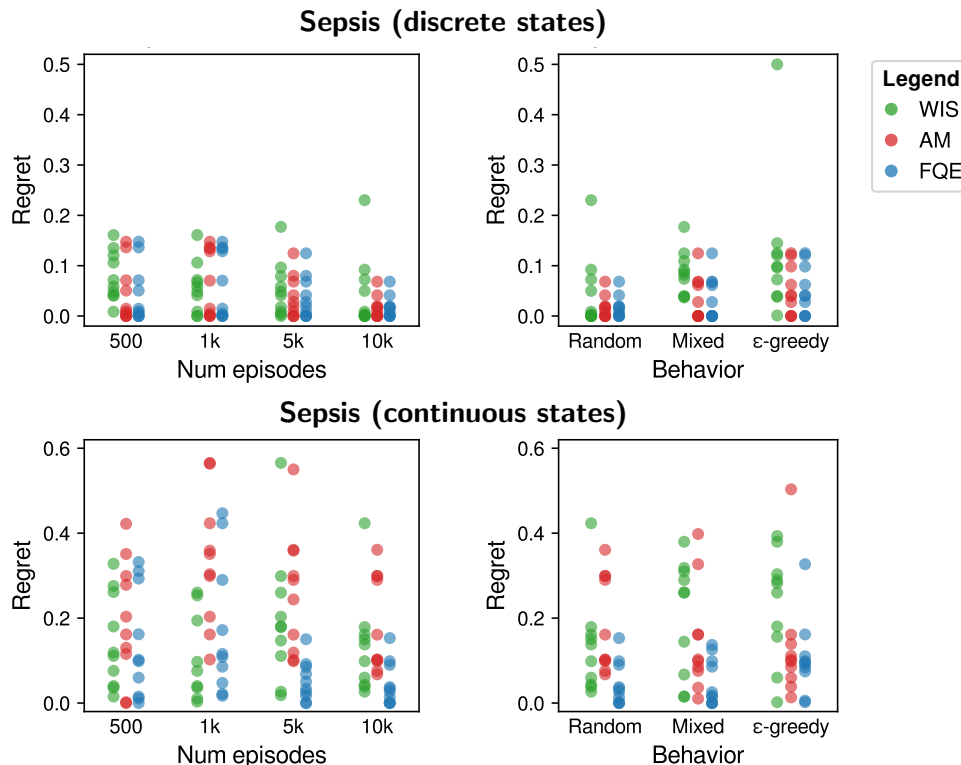


Figure 8: Regret comparison for OPE using validation data with varying sample sizes and varying behaviors. A larger validation set and more exploratory behavior generally results in the lowest regret.

Results and Discussion. For settings with both discrete and continuous states ([Figure 8](#), left column), FQE consistently benefits from larger validation sizes, achieving lower regret. The trend is less apparent for WIS (which generally results in noisier policy value estimates) and for AM (due to model misspecification, especially with function approximation).

D.2.3. TYPE OF BEHAVIOR POLICY AND EXTENT OF EXPLORATION

The behavior policies used to collect validation data can affect which part of the state-action space receives sufficient exploration and the resultant OPE model selection performance. Similar to [Section 5.2.3](#), we consider the random behavior policy, the near-optimal ϵ -greedy behavior policy, and a mixture of both to create the validation sets each containing $m = 10,000$ episodes. The training set and training procedures were kept the same to generate the same candidate set for each run, so that the only difference comes from different validation data.

Results and Discussion. The random behavior policy produces more exploratory data compared to ϵ -greedy, leading to slightly lower regret in general (Figure 8, right column). Surprisingly, even when a mixture of multiple behavior policies is used, WIS is not heavily affected by this policy misspecification (though in the tabular setting, the mixed behavior data produces slightly higher median regret). FQE and AM remain relatively stable across different behavior settings and lead to the lowest regret in most cases with smaller variance.

Appendix E. Additional Related Work

The quality of policies learned by RL methods can depend heavily on hyperparameter choices, especially when combined with deep learning (Henderson et al., 2018). In supervised domains, model selection and hyperparameter tuning has been an essential tool to improve model performance in numerous application settings (Bergstra and Bengio, 2012; Li et al., 2017; Snoek et al., 2012). Yet model selection for offline RL has received relatively less attention, especially in empirical analyses on real-world datasets. Closely related to our setup, Paine et al. (2020) demonstrated that FQE is a useful strategy for hyperparameter selection in offline RL on several benchmark control tasks; Fu et al. (2021) published a benchmark suite for comparing OPE methods on continuous control tasks with Spearman’s correlation as one of evaluation metrics, and also found FQE to be the most effective; Mandel et al. (2014) combined importance sampling with a “temporal” cross-validation technique and applied their approach on an educational game problem. In contrast to these works, we compare multiple OPE estimators both qualitatively and quantitatively to address questions of practical importance for healthcare settings. There are several other notable works in this area, for which we provide a brief overview below. Fard and Pineau (2010) established PAC-Bayesian bounds for offline RL model selection but assumed explicit prior information about environment dynamics or value functions. Farahmand and Szepesvári (2011) proposed BERMIN, which focuses exclusively on Bellman error minimization rather than policy value maximization; however, due to overestimation of value functions during training, these errors/losses are not reflective of the true policy performance (corroborated by our results in Appendix D.1 and by Irpan et al. (2019); Paine et al. (2020) who showed the poor ranking quality of the loss function). Irpan et al. (2019) developed off-policy classification to rank policies, but their approach is limited to goal-directed domains with a sparse binary reward. Lee et al. (2020) proposed a gradient-based hyperparameter optimization procedure but is limited to a single hyperparameter that controls the extent of KL-regularization. Kuzborskij et al. (2021) considered a model selection pipeline similar to ours for contextual bandits and developed a new scoring function, whereas we focus on the sequential setting in offline RL and practical challenges in applying existing OPE methods as scoring functions. Xie and Jiang (2021) studies the problem from a theoretical perspective and proposed BVFT, reducing model selection to pairwise comparisons between candidate policies and provided finite-sample guarantees; in contrast, we consider the more practical aspects of how model selection can (and should) be carried out when applying existing RL/OPE approaches on observational dataset.

Appendix F. Additional Details of the Two-Stage Selection Procedure

In this appendix, we offer some theoretical insight for the expected behavior of the two-stage selection procedure. We first present a general result on the effect of random ranking.

Theorem 1. *Suppose we have a list of C values $[v_{[1]}, \dots, v_{[C]}]$ listed in decreasing order. Given a random permutation $[v_1, \dots, v_C]$ over these C values, the probability that the first A positions contain at least one of the B largest values is:*

$$\Pr(\max\{v_1, \dots, v_A\} \geq v_{[B]}) = \begin{cases} 1 & \text{if } A + B > C \\ 1 - \frac{(C-A)!(C-B)!}{C!(C-A-B)!} & \text{if } A + B \leq C \end{cases}$$

Proof. We consider the following two cases:

- If $A + B > C$, then the B largest values do not fit in the $C - A$ positions because $B > C - A$, and thus the desired probability is 1.
- If $A + B \leq C$, then we can consider the complement, i.e., the probability such that none of the B largest values are in the first A positions. This can be obtained by counting the possibilities of arranging B objects into C positions (for the sample space) and $C - A$ positions (for the event space, such that none of the B objects fall into the first A positions), given by $\frac{(C-A) \cdots (C-A-B+1)}{C \cdots (C-B+1)} = \frac{(C-A)!(C-B)!}{C!(C-A-B)!}$. The desired probability is thus $1 - \frac{(C-A)!(C-B)!}{C!(C-A-B)!}$. \square

Remark. Below we show example values of A , B , C , and the corresponding probability.

C	A	B	Probability
1	1	1	1
K	α	1	α/K
K	1	β	β/K

Remark. We can view this probability as a cdf, with the corresponding pdf being the probability that the maximum value of the first A positions being exactly the B -th largest value $v_{[B]}$ (assuming all values are distinct, wlog). When $A = 1$, we can show that the pdf is $1/C$, regardless of the value of B .

$$\begin{aligned} & \Pr(v_1 = v_{[B]}) \\ &= \Pr(v_1 \geq v_{[B]}) - \Pr(v_1 \geq v_{[B-1]}) \\ &= \left(1 - \frac{(C-1)!(C-(B+1))!}{C!(C-1-(B+1))!}\right) - \left(1 - \frac{(C-1)!(C-B)!}{C!(C-1-B)!}\right) \\ &= \frac{(C-1)!(C-B+1)!}{C!(C-B)!} - \frac{(C-1)!(C-B)!}{C!(C-B-1)!} \\ &= \frac{C-B+1}{C} - \frac{C-B}{C} \\ &= 1/C \end{aligned}$$

This matches with our intuition because a randomly selected value from a list of C values should have $1/C$ probability of being any one of the C values. While this seems trivial, the general formula we derived in [Theorem 1](#) allows us to compute the probability for any values of A and B .

F.1. Performance of Two-Stage Selection

Suppose that a random ranking over candidate policies is used in the first-stage. We can use [Theorem 1](#) to calculate the probability that a good policy is retained in the pruned subset, by setting $C = K$ the number of candidate policies, and $A = \alpha$ the initial subset size. Here, B controls the definition of a “good” policy – where setting $B = \beta$ means a “good” policy is one of the top β policies according to the ground-truth policy values – and this corresponds to the maximum allowable regret.

In [Figure 9](#) (left), we plot this probability against initial subset size α , for different β values. At $\beta = 1$, the probability of achieving zero regret increases linearly as the initial subset size α grows. However, when a larger level of regret is allowed (larger β values), the probability is high even for small α values. This suggests that, for example, pruning the candidate set by half has a high chance of retaining a good policy even if the first-stage OPE has random performance. Since OPE estimators are expected to perform better than random, the above probability provides the lower bound; a (trivial) upper bound can be obtained from an optimal ranking that has probability 1 for any values of A and B , including $A = B = 1$.

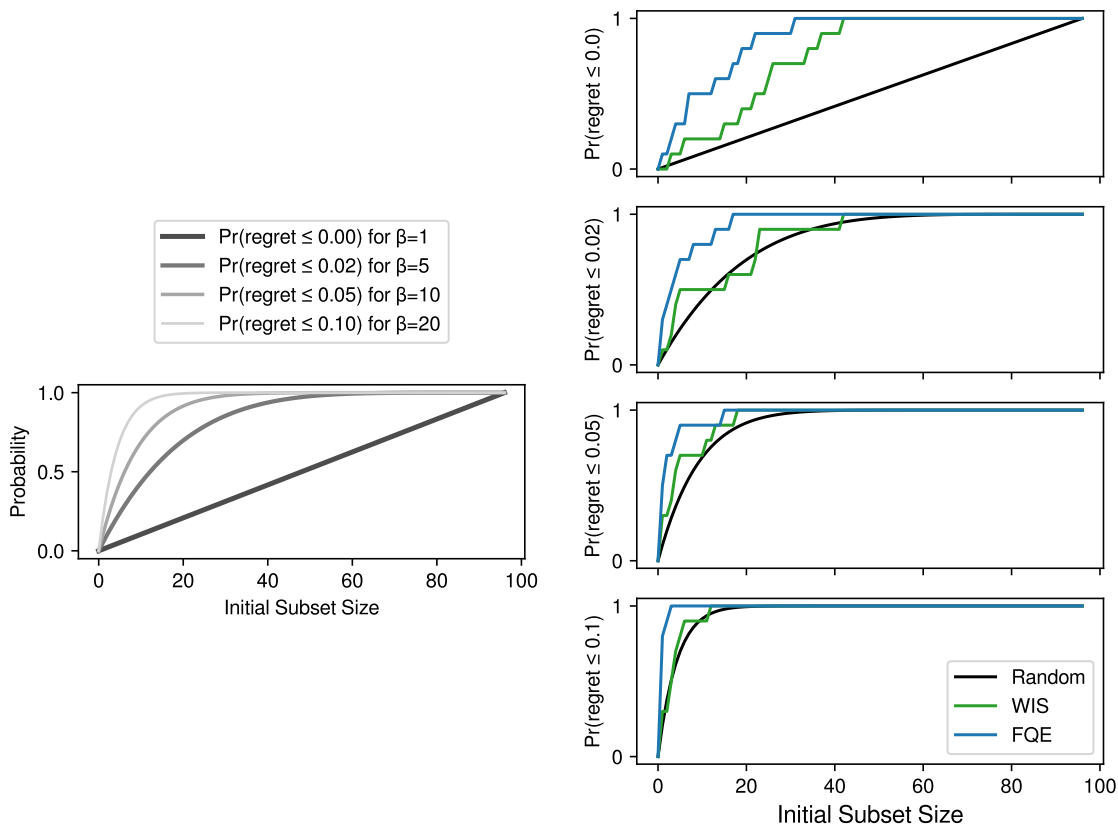


Figure 9: Comparison of cumulative distribution functions of a random OPE, WIS, and FQE, relative to initial subset size α , for the 96 candidate policies on the sepsis simulator task for different β values (which controls the level of allowable regret).

To compare the performance of WIS and FQE against the random lower bound, we computed empirical cdfs from the 10 replication runs for each OPE method, using the main experiment in [Section 5.1](#). In [Figure 9](#) (right), we visualize these cdfs on a set of axes, one for each β value. We see that in general, FQE has a higher chance of retaining a “good” policy than WIS, if used for the first-stage pruning, for all α and β values. Nonetheless, both WIS and FQE have cdfs larger than the random lower bound, suggesting that both OPE estimators indeed perform better than random.

F.2. Further Analyses & Discussion

The analyses above focused on the sole effect of the first-stage pruning, and thus assumed that the second OPE in the two-stage procedure produces a perfect ranking. One could estimate the overall performance by taking the product of probabilities of the two stages, assuming that rankings from the two stages are independent. However, the independence assumption does not hold in reality, because all OPE validation scores are generated from the same validation data. Furthermore, the quality of ranking from an OPE is challenging to quantify as it depends on many factors: the OPE itself, the data, the variance of each value estimate (this can differ per candidate policy), etc. These analyses are outside the scope of this paper and left for future work.