

Learning Action Embeddings for Off-Policy Evaluation

No Author Given

No Institute Given

Abstract. Off-policy evaluation (OPE) methods allow us to compute the expected reward of a policy by using the logged data collected by a different policy. However, when the number of actions is large, or certain actions are under-explored by the logging policy, existing estimators based on inverse-propensity scoring (IPS) can have a high or even infinite variance. Saito and Joachims [12] propose marginalized IPS (MIPS) that uses action *embeddings* instead, which reduces the variance of IPS in large action spaces. MIPS assumes that good action embeddings can be defined by the practitioner, which is difficult to do in many real-world applications. In this work, we explore *learning* action embeddings from logged data. In particular, we use intermediate outputs of a trained reward model to define action embeddings for MIPS. This approach extends MIPS to more applications, and in our experiments improves upon MIPS with pre-defined embeddings, as well as standard baselines, both on synthetic and real-world data. Our method does not make assumptions about the reward model class, and supports using additional action information to further improve the estimates. The proposed approach presents an appealing alternative to DR for combining the low variance of DM with the low bias of IPS.

Keywords: off-policy evaluation · multi-armed bandits · large action space · representation learning · recommender systems

1 Introduction

The *multi-armed bandit* (MAB) framework is commonly used to model the interaction between recommender systems or search engines with their users [20]. In MAB, an *agent* performs *actions* and receives *rewards*, where each reward is sampled from an unknown reward distribution conditioned on the action. The goal is to learn a *policy* for the agent that maximizes the cumulative reward over multiple iterations. In a *contextual* MAB, the reward distribution is also conditioned on a *context* that is observed by the agent, and is used as an additional input to the policy function.

Bandit policies are commonly learned *online*, where we run a policy in production and update it iteratively using the observed rewards [20]. However, deploying untested policies to production is risky. An alternative is to leverage the logged data from historical customer interactions to learn a policy *offline*.

The fundamental problem of offline learning is *off-policy evaluation* (OPE), where we try to estimate the expected reward of a new policy without deploying it. The disadvantage of OPE is that the computed value is only an *estimate* of the true value of the policy, and the success of offline learning depends on how accurate this estimate is. When the policy used to gather logged data has a low probability of choosing a subset of actions, existing estimators based on *inverse propensity scoring* [IPS, 9] can be inaccurate [10]. This problem is especially potent when the number of actions is large, as is often the case in recommender systems [18].

Embedding contexts into a lower-dimensional space can improve the off-policy estimates when dealing with a large number of contexts [20]. Similarly, Saito and Joachims [12] propose to embed the actions, and to use the embeddings in *marginalized IPS* (MIPS). Saito and Joachims use additional action information to define the embeddings. For example, in fashion recommendation, where an action is a particular fashion item to recommend, they use the price, the brand and the category of an item as embedding dimensions. For many problems of interest defining action embeddings by hand will be difficult, and/or will require expert knowledge.

Motivated by the two considerations above, we propose methods for *learning* or *fine-tuning* the action embeddings, aiming to improve the performance of MIPS. Our key hypothesis is that an action embedding that is useful for reward prediction will also be useful for MIPS. In particular, we propose to use the intermediate outputs of a trained reward model as action embeddings for MIPS, instead of using them for direct reward prediction as in DM.

We study this approach on both synthetic and real-world datasets, demonstrating that it consistently outperforms MIPS with pre-defined action embeddings, as well as common baselines like standard IPS, *direct method* (DM), and *doubly robust* estimation [DR, 2]. We demonstrate that the method is not sensitive to reward model misspecification: a linear regression reward model produces useful action embeddings even when the true reward function itself is non-linear. We propose methods for utilizing additional action information when it is available, and show that learning/fine-tuning action embeddings is especially useful when the additional action information is high-dimensional, or when the information affecting the reward is incomplete. Finally, we provide a theoretical analysis of our method, showing that it can be interpreted as kernelized DM.

2 Background and related work

The interaction of a recommender system with the user starts with the system receiving the context information (e.g. user’s purchase history) $x \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$ drawn i.i.d. from an unknown distribution $p(x)$. The system then chooses an action (e.g. a ranked set of products) $a \sim \pi(a | x)$ from the set of available actions \mathcal{A} according to the system’s policy π . Finally, the user interacts with the displayed products via clicks or purchases, and the system receives a reward $r \in [r_{min}, r_{max}]$ sampled from a reward distribution $p(r | x, a)$. This process repeats for n rounds, where in round t the system observes a context x_t , draws

an action a_t according to $\pi(a | x_t)$, and observes a reward $r_t \sim p(r | x_t, a_t)$. All interactions are recorded in a *logged dataset* $\mathcal{D} = \{(x_t, a_t, r_t)\}_{t=1}^n$. The policy π_0 that collects \mathcal{D} is called a *logging policy*. The goal of *off-policy evaluation* [OPE, 2, 19, 15, 14, 7, 12] is to develop \hat{V} , an estimator of the policy *value* $V(\pi) = \mathbb{E}_{p(x)\pi(a|x)p(r|x,a)}[r]$ such that $\hat{V}(\pi) \approx V(\pi)$.

Inverse propensity scoring [IPS, 9] is a commonly used estimator that works by re-weighting the rewards in \mathcal{D} based on how likely the corresponding actions are to be selected under the target policy:

$$\hat{V}_{\text{IPS}}(\pi) = \frac{1}{n} \sum_{t=1}^n \frac{\pi(a_t | x_t)}{\pi_0(a_t | x_t)} r_t. \quad (1)$$

Many state-of-the-art OPE methods are based on IPS [2, 19, 15, 14, 7]. A key assumption of IPS, however, is that the logging policy π_0 and the target policy π have *common support*, i.e. $\pi(a | x) > 0 \implies \pi_0(a | x) > 0$ [12]. When dealing with large action spaces, it can be difficult to select a logging policy that puts non-zero probability on *all* actions, yet does not degrade the customer experience. Even if the two policies have common support, a low probability of selecting an action under the logging policy results in a large weight $\pi(a_t|x_t)/\pi_0(a_t|x_t)$, which increases the variance of the IPS estimate. To reduce the variance of IPS, we can either clip the importance weights according to a tunable clipping parameter [16] or self-normalize them [SNIPS, 17], but this biases the estimator. To deal with continuous space, Kallus and Zhou [6] use a kernel function to calculate propensities for IPS. We also use a similar technique to estimate the propensities of learned embeddings, and we show that the propensity estimation used in MIPS can be interpreted as kernel regression in Section 3.3.

In their work Saito and Joachims [12] instead assume that useful *action embeddings* are available to the estimator, and that we can pool information across similar actions to improve the estimates. Suppose we have an action embedding $e \in \mathcal{E} \subseteq \mathbb{R}^{d_e}$ for each action a . Saito and Joachims show that if the logging policy π_0 has common *embedding support*, i.e. $p(e | \pi, x) > 0 \implies p(e | \pi_0, x) > 0$, and the action a has no *direct effect* on the reward r , i.e. $a \perp r | x, e$, then the proposed *marginalized IPS* (MIPS) estimator

$$\hat{V}_{\text{MIPS}}(\pi) = \frac{1}{n} \sum_{t=1}^n \frac{p(e_t | \pi, x_t)}{p(e_t | \pi_0, x_t)} r_t \quad (2)$$

is unbiased. Experiments by Saito and Joachims demonstrate that in practice MIPS has a lower variance than IPS, and produces estimates with a lower MSE. Our goal in this work is to develop a method for *learning* the action embeddings from logged data to further improve MIPS, and to make it applicable to problems where pre-defined embeddings are not available.

The *direct method* (DM) learns the model of the expected reward for the context-action pair, and uses it to estimate the expected reward of a policy. DM has a low variance, but suffers from a high bias when the model is misspecified. *Doubly robust* estimator [DR, 2, 4, 19] combines DM and IPS: it uses the learned

model of the reward as a control variate. In other words, it computes a non-parametric estimation on the reward model residuals. If the expected reward is correctly specified, DR can achieve a lower variance than IPS, and a lower bias than DM. Several extensions of DR have been proposed [19, 14], but, to the best of our knowledge, none of them aim to improve the performance of DR in large action spaces.

Two methods for off-policy evaluation with large action spaces have been developed in parallel with our work [8, 13]. Peng et al. [8] propose to cluster similar actions, and replace individual action propensity scores with those of action clusters. Saito et al. [13] train a two-step regression model, hypothesizing that similar actions share a cluster effect on the reward, but also have their own residual effect.

3 Methods

In this work we propose to use the reward signal to learn the embeddings of actions for the MIPS estimator, which would estimate the propensities of these embeddings, and use them instead of the original action propensities to re-weight the observed rewards.

3.1 Motivation

We first formalize the main assumption that motivates the proposed method.

Assumption 1 (Action Embedding Space) *For every set of actions \mathcal{A} , there exists a lower-dimensional embedding space $\mathcal{E} \subseteq \mathbb{R}^{d_{\mathcal{E}}}$ so that every action $a \in \mathcal{A}$ can be mapped to an embedding $e \in \mathcal{E}$ while $p(r | a, x) \approx p(r | e, x)$ holds for every context x . We denote $d_{\mathcal{E}} < |\mathcal{A}|$ to be the number of dimensions of embedding space \mathcal{E} .*

This assumption holds when $d_{\mathcal{E}} = |\mathcal{A}|$ as the actions can be mapped as one-hot encoded representations. But we hypothesize that in practice some actions are “similar”, and that the action space can be represented more efficiently. For example, if certain types of customers mostly clicked on shoes in the logged data, we expect to observe a high reward for a shoe product, even if we have never recommended this particular product for this type of customer before.

Based on Assumption 1, the environment can be represented as a graphical model $a \rightarrow e \rightarrow r \leftarrow x$. This can be learned the same way as learning a reward model for DM, but instead of using the model for reward prediction, we use the intermediate model outputs as action embeddings. As model misspecification is one of the biggest issues with DM [4], it is often difficult to predict the reward end-to-end. But even when the true reward function is complex, and the model class is not rich enough to recover it, we hypothesize that the model will learn a useful $a \rightarrow e$ mapping that can be used in MIPS as defined in Equation (2).

Table 1: Mean squared error of the estimators on a synthetic experiment while varying the number of actions. Reporting mean and standard error averaged over 50 runs.

$ \mathcal{A} $	IPS	DM	LEARNED MIPS
50	0.63 ± 0.05	0.64 ± 0.05	0.56 ± 0.05
100	0.68 ± 0.04	2.13 ± 0.11	0.58 ± 0.05
200	0.61 ± 0.04	7.16 ± 0.19	0.55 ± 0.04
500	0.71 ± 0.05	27.9 ± 0.32	0.56 ± 0.05
1000	0.70 ± 0.04	62.3 ± 0.31	0.55 ± 0.03

To test our hypothesis, we design the following toy example¹. Let $\mathcal{D} = \{(x_t, a_t, r_t)\}_{t=1}^n$ for $n = 1000$ be the logged dataset, where a context $x \in [-1, 1]^{d_{\mathcal{X}}}$, $d_{\mathcal{X}} = 5$, is drawn from the standard normal distribution, action $a \sim \pi_0$ is chosen from set \mathcal{A} by the unconditional logging policy $\pi_0(a) = \nu_a / \sum_{a' \in \mathcal{A}} \nu_{a'}$, where $\nu_a \sim \text{Exp}(1)$ is drawn from an exponential distribution, and the reward is generated by $r = f_a(x) + \eta$, where $f_a(x) = 1 / (1 + e^{-(x^\top \theta_a + \mu_a)})$ is a logistic model for action a with randomly initialized parameters $\theta_a \sim \mathcal{N}(0, 1)^{d_{\mathcal{X}}}$, $\mu_a \sim \mathcal{N}(0, 1)$, and $\eta \sim \mathcal{N}(0, 0.1)$ is random noise. We use the uniform random target policy. For DM, we fit a linear regression $\hat{f}_a(x)$ to estimate r . Of course, linear regression is not expressive enough to fit the non-linear f_a , but it learns the correct *order* of the contexts, meaning $\forall x_1, x_2 \in \mathcal{X}, \hat{f}_a(x_1) < \hat{f}_a(x_2) \Leftrightarrow f_a(x_1) < f_a(x_2)$. We take these learned embeddings and use them in the MIPS estimator [11] and refer to this method as *Learned MIPS*. We vary the number of actions $|\mathcal{A}| \in \{50, 100, 200, 500, 1000\}$, while keeping the size of the logged data fixed. For each $|\mathcal{A}|$ we synthesize 50 reward functions, and for each reward function we generate 15 logged datasets to evaluate the methods. The results are summarized in Table 1.

As the number of samples per action decreases, the learned embeddings become increasingly inaccurate for *direct* reward prediction (see the increasing error of DM in Table 1), but they still reflect the structure of the problem, and we can build an accurate model-free estimator on top of them (Learned MIPS outperforms both DM and IPS). This simple example demonstrates the potential of the proposed method as a novel way to combine the strengths of model-based and model-free off-policy estimation methods.

3.2 Algorithm

As shown in the previous section, we can learn action embeddings by fitting a separate linear regression model $\hat{f}_a(x)$ for each action a . In general, as in DM, we can use any model class. For example, we can fit a deep neural network, and use the intermediate outputs as action embeddings. In this work, we use a linear model visualized in Figure 1. In our preliminary experiments, more

¹ Code to reproduce all experiments will be made public following publication.




Fig. 1: The model we use to learn the action embeddings. We use a simple linear model in this work, but any model class can be used. Input a can be a one-hot encoded representation of the action identity, a pre-defined action embedding, or a concatenated vector of both.

complex neural network models demonstrated comparable or worse performance. We hypothesize that this is due to the high expressivity (and hence high variance) of neural networks, which is counter to our goal of reducing estimator variance. We leave further empirical study and theoretical analysis of this phenomenon for future work.

The model takes the action identity a , and computes a reward estimate \hat{r} as a dot product between the action embedding and the context vector x . We train the model to minimize the MSE between \hat{r} and the reward r observed in the logged data. The learned embedding for action a is the output of the linear layer. As we show in Section 4, we can also include additional action information (e.g. content-based features of the corresponding product), which can further improve the performance of the estimator.

Having learned the action embeddings, we follow MIPS [12] and fit a logistic regression model to estimate $p(a | e)$. This estimate is used to compute the embedding propensities ($p(e | \pi, x)$ and $p(e | \pi_0, x)$) in Equation (2), noting that $p(e | \pi, x) = p(a | e)\pi(a | x)$. We discuss this step in more detail in the next section, where we show a connection between our method and kernelized DM.

3.3 Connection to DM

Learned MIPS resembles DM in that both methods solve the OPE task by first training a reward predictor. The two methods differ, however, in how they apply said predictor. In this section, we show a direct connection between learned MIPS and DM. In particular, we show that we can interpret the proposed method as DM that uses *kernel regression* with a learned feature embedding as the reward predictor.

To simplify the notation, we consider a non-contextual estimation task: given logged data $\mathcal{D} = \{a_t, r_t\}_{t=1}^n$ gathered with a policy $\pi_0(a)$, we want to estimate

the value of a policy $\pi(a)$. In this setting, the DM estimate is simply

$$\hat{V}_{\text{DM}}(\pi) = \sum_{a \in \mathcal{A}} \pi(a) \hat{r}(a), \quad (3)$$

where $\hat{r}(a)$ is the model of expected reward for action a that we learn using \mathcal{D} as the training data. We now define $\tilde{\mathcal{D}} = \{a_t, e_t, r_t\}_{t=1}^n$, where $e_t \sim p(e | a_t)$ is the corresponding action embedding. The MIPS estimate is then

$$\hat{V}_{\text{MIPS}}(\pi) = \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{p(a|e_t)} \left[\frac{\pi(a)}{\pi_0(a)} \right] r_t = \frac{1}{n} \sum_{t=1}^n \sum_{a \in \mathcal{A}} \frac{\pi(a)}{\pi_0(a)} p(a | e_t) r_t \quad (4)$$

$$= \sum_{a \in \mathcal{A}} \pi(a) \frac{1}{n} \sum_{t=1}^n \frac{p(a | e_t)}{\pi_0(a)} r_t = \sum_{a \in \mathcal{A}} \pi(a) \tilde{r}(a), \quad (5)$$

where $p(a | e_t)$ is estimated as in Equation (7). In other words, the MIPS estimate matches the DM estimate with a different expected reward model:

$$\tilde{r}(a) = \frac{1}{n} \sum_{t=1}^n \frac{p(a | e_t)}{\pi_0(a)} r_t. \quad (6)$$

We now try to understand the nature of \tilde{r} . Intuitively, $p(a | e)$ will be higher when the embedding of a is “closer” to e according to some distance measure. In particular, let $e \sim \mathcal{N}(f(a), \sigma^2 I)$ be the embedding distribution, where $f(a)$ is the learned deterministic embedding. As in MIPS, we estimate $p(a | e)$ by fitting a probabilistic classifier to the dataset $\{a_t, e_t\}_{t=1}^n$. Saito and Joachims use logistic regression, but here we consider *linear discriminant analysis* [LDA, 5, Section 4.3], a comparable model. LDA fits a normal distribution to each class, hence we expect it to recover $\mathcal{N}(f(a), \sigma^2 I)$. We also expect it to recover $\pi_0(a)$ as the class prior, because we sampled $a_t \sim \pi_0(a)$ to produce the training data $\{a_t, e_t\}_{t=1}^n$. The predictive probability of LDA is then

$$p(a | e) \propto \mathcal{N}(e; f(a), \sigma^2 I) \pi_0(a). \quad (7)$$

Combining Equations (6) and (7) we get

$$\tilde{r}(a) \propto \frac{1}{n} \sum_{t=1}^n \mathcal{N}(e_t; f(a), \sigma^2 I) r_t, \quad (8)$$

where

$$\mathcal{N}(e_t; f(a), \sigma^2 I) \propto \exp \left(-\frac{1}{2\sigma^2} (e_t - f(a))^\top (e_t - f(a)) \right). \quad (9)$$

This confirms our intuition: if we use LDA as the classifier, the expected reward $\tilde{r}(a)$ in Equation (5) is a weighted combination of all rewards in the logged data, where samples with e_t “closer” to the action embedding $f(a)$ have a higher

weight. For the normal embedding distribution, the weight is determined by a squared Euclidean distance between vectors e_t and $f(a)$, but other distributions will induce different distance functions. In fact, Equation (8) is equivalent to *kernel regression* [1, Chapter 10] in a learned embedding space determined by f , with a kernel determined by the embedding distribution.

While the above holds when using LDA as the classifier, in our experiments we follow Saito and Joachims and use the logistic regression classifier. Both of these methods fit linear decision boundaries and give similar results in practice (3; 5, Chapter 4), which means that our interpretation likely holds for logistic regression-based MIPS as well. As DM, our method does not provide theoretical guarantees on the bias/variance for a general reward model class, but, as we show in the next section, demonstrates good performance in practice.

4 Empirical evaluation

In this section we extend our proof-of-concept experiments in Section 3 to a more complex synthetic dataset, and conduct experiments on a real-world fashion recommendation dataset. In the synthetic experiments, we evaluate two classes of methods based on whether the method uses additional action information or not. We find that our methods outperform all the standard baselines and, in some cases, improve upon the true embeddings used to generate the reward.

4.1 Synthetic experiments

For our empirical evaluation, we follow the setup by Saito and Joachims [12]. We generate synthetic data by first sampling $d_{\mathcal{X}}$ -dimensional context vectors x from the standard normal distribution. We also sample $d_{\mathcal{E}}$ -dimensional categorical action embeddings $e \in \mathcal{E}$ from the following distribution

$$p(e | a) = \prod_{k=1}^{d_{\mathcal{E}}} \frac{\exp \alpha_{a,e_k}}{\sum_{e' \in \mathcal{E}_k} \exp \alpha_{a,e'}}, \quad (10)$$

where $\{\alpha_{a,e_k}\}$ is a set of parameters sampled independently from the standard normal distribution. Each dimension of \mathcal{E} has a cardinality of 10. We then synthesize the expected reward as

$$q(x, e) = \sum_{k=1}^{d_{\mathcal{E}}} \eta \cdot (x^{\top} M x_{e_k} + \theta_x^{\top} x + \theta_e^{\top} x_{e_k}), \quad (11)$$

where M is the parameter matrix and θ_x, θ_e are parameter vectors sampled uniformly in $[-1, 1]$. x_{e_k} is a parameter vector corresponding to the k -th dimension of the action embedding and is sampled from the standard normal distribution. Parameter η_k specifies the importance of the k -th dimension and is drawn from a Dirichlet distribution $\eta \sim \text{Dir}(\alpha)$ with its parameters $\alpha = (\alpha_1, \dots, \alpha_{d_{\mathcal{E}}})$ sampled from a uniform distribution $[0, 1]$. With this setup, the action embeddings are

fully informative, and the action identity has no *direct effect* on the reward. On the other hand, if action embeddings have some missing dimensions, they do not fully capture the action’s effect on the reward. In this case, we hypothesize that we can learn better embeddings, and we test this in one of the experiments.

We also follow Saito and Joachims [12] in how we synthesize the logging policy π_0 . We apply the softmax function to $q(x, a) = \mathbb{E}_{p(e|a)}[q(x, e)]$ as

$$\pi_0(a | x) = \frac{\exp \beta \cdot q(x, a)}{\sum_{a' \in \mathcal{A}} \exp \beta \cdot q(x, a')}, \quad (12)$$

where β is the parameter that controls the entropy and the optimality of the logging policy. A large positive value leads to an optimal, near-deterministic logging policy, while 0 leads to a uniform-random policy. We define the target policy π as an ε -greedy policy

$$\pi(a | x) = (1 - \varepsilon) \cdot \mathbb{1} \left\{ a = \arg \max_{a' \in \mathcal{A}} q(x, a') \right\} + \frac{\varepsilon}{|\mathcal{A}|}, \quad (13)$$

where $\varepsilon \in [0, 1]$ controls the exploration rate of the target policy. Similar to Saito and Joachims [12], we set $\beta = -1$ and $\varepsilon = 0.05$ in our experiments.

To generate a batch of logged data, we follow an iterative process to generate n samples $\mathcal{D} = \{x_t, a_t, e_t, r_t\}_{t=1}^n$. In each iteration, we sample a context x from the standard normal distribution ($d_{\mathcal{X}} = 10$), and a discrete action a from π_0 according to Equation (12). Given an action a , we then sample a categorical action embedding e using Equation (10). Finally, we sample the reward from a normal distribution with mean $q(x, e)$ defined in Equation (11) and standard deviation $\sigma = 2.5$ [12].

We assume three variants of our method, based on which action representation is passed to the model:

- *Learned MIPS OneHot* only sees one-hot encoded action identities.
- *Learned MIPS FineTune* only sees pre-defined action embeddings.
- *Learned MIPS Combined* sees both of the above concatenated into a single vector.

We compare our methods to all common estimators, including DM, IPS, and DR. To gauge the quality of the learned embeddings, we specifically compare to MIPS with pre-defined embeddings, and *MIPS (true)*, which uses the true propensity scores of the pre-defined embeddings, which are not available in practice. For all experiments we report the MSE averaged over 100 different subsampled datasets.

Can we improve upon standard baselines without using pre-defined embeddings?

We evaluate our performance against all standard estimators in two sets of experiments with data generated according to the procedure described in Section 4.1. All pre-defined embedding dimensions are preserved, hence these embeddings are *fully-informative*, i.e. $q(x, e) = q(x, e, a)$. We fix the number of pre-defined embedding dimensions to $d_{\mathcal{E}} = 3$. In the first experiment, we vary the number of actions




Fig. 2: Synthetic experiments varying the number of actions and training samples. To better distinguish comparable methods, those marked with ∇ only use action identities, and those marked with \times also use pre-defined action embeddings. *Learned MIPS OneHot* outperforms all standard baselines. When we have enough data for every action, it performs just as well as IPS. As the variance grows with fewer samples per action, its error approaches the one of DM using the same model. The pre-defined embeddings ($d_{\mathcal{E}} = 3$) have low bias and variance; hence our methods can not improve upon them.

from 10 to 2,000 with a fixed sample size of $n = 20,000$. In the second experiment, we vary the number of samples in the training data $n = [800, 1600, \dots, 102400]$ and leave the number of actions fixed at $|\mathcal{A}| = 100$. The results of these two experiments are presented in Figure 2.

Learned MIPS OneHot outperforms all other methods using action identities. We can observe the difference between the two methods that use the same model, *Learned MIPS OneHot* and *DM*. Using learned embedding propensities in MIPS results in a smaller bias compared to using them in direct reward prediction. The error of *Learned MIPS* keeps shrinking with more data. While IPS and DR do not pool information across actions, MIPS is using the embedding propensities influenced by every action, resulting in a better bias-variance trade-off. When the number of actions is small enough, and we have enough data, IPS estimates already have a low variance, and we do not observe any significant improvements.

We also report the performance of the methods that use pre-defined embeddings. Perhaps unsurprisingly, the learned embeddings perform worse than pre-defined embeddings, and even *Learned MIPS FineTune* does not improve over them. This is because the pre-defined embeddings compress the information about the action well. In further experiments, we study more realistic benchmarks, where embeddings are higher-dimensional or do not fully capture the reward. Both of these cases are common in practice, for example, when using pre-defined embeddings extracted from LLMs.

Can we improve upon the pre-defined embeddings? In the previous experiment, fine-tuned embeddings did not improve upon the pre-defined embeddings. The reason is that the pre-defined embeddings have low bias and low variance. This is rarely the case in practice. In this experiment, we show that we can learn better

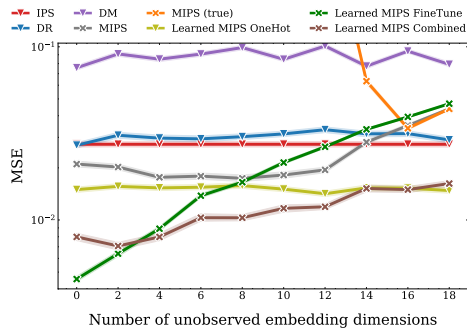


Fig. 3: Synthetic experiments varying the number of unobserved dimensions. As we progressively hide more dimensions, pre-defined embeddings get more biased and methods using them get less accurate. Combining *Learned MIPS OneHot* and *Learned MIPS FineTune* yields the most robust results when the bias of pre-defined embeddings is unknown.

embeddings when the pre-defined embeddings have a high number of embedding dimensions or when they do not capture the full effect on the reward. We evaluate how the bias of the pre-defined embeddings influences our methods and whether it is better to learn the embeddings from action identities or fine-tune them from pre-defined embeddings in this setting. We progressively add more bias similarly to Saito and Joachims [12]. We fix the number of dimensions in the action embedding to $d_{\mathcal{E}} = 20$, and after the reward is generated, we hide a certain number of dimensions so the estimators do not have access to fully-informative embeddings. The results are shown in Figure 3. As the embeddings worsen, our methods that learn or fine-tune them outperform MIPS. We can see the bias-variance trade-off between *Learned MIPS OneHot* and *Learned MIPS FineTune*. When the number of unobserved dimensions is small, the variance reduction gained from pre-defined embeddings is greater than induced bias. As we increase the number of unobserved dimensions, the bias of pre-defined embeddings can get even higher than the variance of IPS. In practice, we often do not know how biased the pre-defined embeddings are. Therefore, we introduce *Learned MIPS Combined* and observe performance improvements when using this method for the majority of the values of the unobserved dimensions.

4.2 Real-world data

We follow the experimental setup of Saito and Joachims [12] to evaluate how the estimators perform on a real-world bandit dataset. In particular, we use a fashion recommendation dataset that was collected on a large-scale fashion e-commerce platform, and comes in the form of $\mathcal{D} = \{(x_t, a_t, e_t, r_t)\}_{t=1}^n$ tuples. The dataset was collected during an A/B test under two policies: a uniform policy π_0 and a Thompson sampling policy π . The size of the action space is $|\mathcal{A}| = 240$, and every action (a fashion item) comes with a pre-defined 4-dimensional action embedding e , which includes the item’s price, brand, and hierarchical category information.

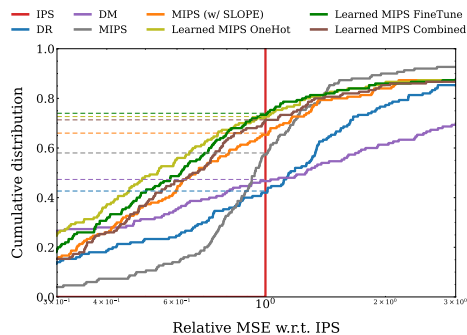


Fig. 4: CDF of relative MSEs w.r.t IPS on the real-world dataset. The intersection of a method’s curve with the IPS curve tells us the proportion of experiments in which the method performs better than IPS.

We use $n = 10000$ sub-sampled observations from the “ALL” campaign in the dataset to be directly comparable to the prior work [12, 8].

In line with Saito and Joachims [12], we repeat the experiment over 150 bootstrap samples, normalize the estimators’ MSE relative to the MSE of IPS (i.e. divide the method’s MSEs for a given data sample by the MSE of IPS for the same sample), and compute the cumulative distribution function (CDF) of these MSEs. We report the results in Figure 4. We compare against another competitive baseline *MIPS(w/ SLOPE)* that greedily improves the performance by dropping embedding dimensions [12]. *Learned MIPS* methods outperform IPS in about 75% of runs, more often than any other baseline. The next best-performing method is MIPS (w/ SLOPE), which outperforms IPS in about 65% of runs. For DM, we use the same model as in *Learned MIPS OneHot*. These results support our hypothesis that even though DM fails to model the reward, the embeddings it learns in the process can still be useful for model-free estimation.

5 Conclusion

In this work, we propose methods that assume a structured action space in the contextual bandit setting, and reduce the variance of model-free off-policy estimators by *learning* action embeddings from the reward signal. We extend MIPS to settings where pre-defined embeddings are not available, or are difficult to define. At the same time, we show that the method can improve upon the pre-defined embeddings even when they are available. The proposed method outperforms all standard baselines on synthetic and real-world fashion recommendation data, even if the reward model itself is inaccurate. An interesting future direction is to study the embeddings learned by more complex classes of reward models, such as neural networks. We would also like to extend our methods to the learning-to-rank setting, as list-level embeddings can dramatically reduce the combinatorial complexity of the problem.

Bibliography

- [1] Dhrymes, P.J.: Topics in advanced econometrics: Probability foundations, vol. 1. Springer Science & Business Media (1989)
- [2] Dudík, M., Erhan, D., Langford, J., Li, L.: Doubly robust policy evaluation and optimization. *Statistical Science* **29**(4), 485–511 (2014)
- [3] Efron, B.: The efficiency of logistic regression compared to normal discriminant analysis. *Journal of the American Statistical Association* **70**(352), 892–898 (1975)
- [4] Farajtabar, M., Chow, Y., Ghavamzadeh, M.: More robust doubly robust off-policy evaluation. In: *International Conference on Machine Learning*, pp. 1447–1456, PMLR (2018)
- [5] Hastie, T., Tibshirani, R., Friedman, J.H., Friedman, J.H.: The elements of statistical learning: data mining, inference, and prediction, vol. 2. Springer (2009)
- [6] Kallus, N., Zhou, A.: Policy evaluation and optimization with continuous treatments. In: *International conference on artificial intelligence and statistics*, pp. 1243–1251, PMLR (2018)
- [7] Metelli, A.M., Russo, A., Restelli, M.: Subgaussian and differentiable importance sampling for off-policy evaluation and learning. *Advances in Neural Information Processing Systems* **34**, 8119–8132 (2021)
- [8] Peng, J., Zou, H., Liu, J., Li, S., Jiang, Y., Pei, J., Cui, P.: Offline Policy Evaluation in Large Action Spaces via Outcome-Oriented Action Grouping. In: *Proceedings of the ACM Web Conference 2023*, pp. 1220–1230, WWW '23, Association for Computing Machinery, New York, NY, USA (Apr 2023), ISBN 978-1-4503-9416-1, <https://doi.org/10.1145/3543507.3583448>
- [9] Robins, J.M., Rotnitzky, A., Zhao, L.P.: Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* **89**(427), 846–866 (1994), <https://doi.org/10.1080/01621459.1994.10476818>, URL <https://doi.org/10.1080/01621459.1994.10476818>
- [10] Sachdeva, N., Su, Y., Joachims, T.: Off-policy bandits with deficient support. In: Gupta, R., Liu, Y., Tang, J., Prakash, B.A. (eds.) *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Virtual Event, CA, USA, August 23-27, 2020, pp. 965–975, ACM (2020), <https://doi.org/10.1145/3394486.3403139>, URL <https://doi.org/10.1145/3394486.3403139>
- [11] Saito, Y., Aihara, S., Matsutani, M., Narita, Y.: Open bandit dataset and pipeline: Towards realistic and reproducible off-policy evaluation. arXiv preprint arXiv:2008.07146 (2020)
- [12] Saito, Y., Joachims, T.: Off-policy evaluation for large action spaces via embeddings. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) *Proceedings of the 39th International Conference on Machine Learning*, *Proceedings of Machine Learning Research*, vol. 162, pp.

- 19089–19122, PMLR (17–23 Jul 2022), URL <https://proceedings.mlr.press/v162/saito22a.html>
- [13] Saito, Y., Ren, Q., Joachims, T.: Off-Policy Evaluation for Large Action Spaces via Conjunct Effect Modeling. In: Proceedings of the 40th International Conference on Machine Learning, pp. 29734–29759, PMLR (Jul 2023), ISSN 2640-3498
 - [14] Su, Y., Dimakopoulou, M., Krishnamurthy, A., Dudík, M.: Doubly robust off-policy evaluation with shrinkage. In: International Conference on Machine Learning, pp. 9167–9176, PMLR (2020)
 - [15] Su, Y., Wang, L., Santacatterina, M., Joachims, T.: Cab: Continuous adaptive blending for policy evaluation and learning. In: International Conference on Machine Learning, pp. 6005–6014, PMLR (2019)
 - [16] Swaminathan, A., Joachims, T.: Counterfactual risk minimization: Learning from logged bandit feedback. In: International Conference on Machine Learning, pp. 814–823, PMLR (2015)
 - [17] Swaminathan, A., Joachims, T.: The self-normalized estimator for counterfactual learning. *advances in neural information processing systems* **28** (2015)
 - [18] Swaminathan, A., Krishnamurthy, A., Agarwal, A., Dudik, M., Langford, J., Jose, D., Zitouni, I.: Off-policy evaluation for slate recommendation. *Advances in Neural Information Processing Systems* **30** (2017)
 - [19] Wang, Y.X., Agarwal, A., Dudik, M.: Optimal and adaptive off-policy evaluation in contextual bandits. In: International Conference on Machine Learning, pp. 3589–3597, PMLR (2017)
 - [20] Zhou, L.: A survey on contextual multi-armed bandits. arXiv preprint arXiv:1508.03326 (2015)