

SPOT: A framework for selection of prototypes using optimal transport

Karthik S. Gurumoorthy ^{*1}, Pratik Jawanpuria^{*2}, and Bamdev Mishra^{*2}

¹ India Machine Learning, Amazon, Bangalore, India

gurumoor@amazon.com

² Microsoft, Hyderabad, India

{pratik.jawanpuria, bamdevm}@microsoft.com

Abstract. In this work, we develop an optimal transport (OT) based framework to select informative prototypical examples that best represent a given target dataset. Summarizing a given target dataset via representative examples is an important problem in several machine learning applications where human understanding of the learning models and underlying data distribution is essential for decision making. We model the prototype selection problem as learning a sparse (empirical) probability distribution having the minimum OT distance from the target distribution. The learned probability measure supported on the chosen prototypes directly corresponds to their importance in representing the target data. We show that our objective function enjoys a key property of submodularity and propose an efficient greedy method that is both computationally fast and possess deterministic approximation guarantees. Empirical results on several real world benchmarks illustrate the efficacy of our approach.

Keywords: Targeted prototype selection · Data subset selection · Optimal transport · Submodularity · Parallelizable greedy method

1 Introduction

Extracting informative and influential samples that best represent the underlying data-distribution is a fundamental problem in machine learning [53,2,31,33,55]. As sizes of datasets have grown, summarizing a dataset with a collection of representative samples from it is of increasing importance to data scientists and domain-specialists [3]. Prototypical samples offer interpretative value in every sphere of humans decision making where machine learning models have become integral such as healthcare [5], information technology [26], and entertainment [42], to name a few. In addition, extracting such compact synopses play a pivotal tool in depicting the scope of a dataset, in detecting outliers [30], and for compressing and manipulating data distributions [43]. Going across domains to identify representative examples from a source set that explains a different target set have recently been applied in model agnostic Positive-Unlabeled (PU) learning [13]. Existing works [52] have also studied the generalization properties of machine learning models trained on a prototypical subset of a large dataset.

* equal contribution

Works such as [2,8,54,52] consider selecting representative elements (henceforth also referred to as *prototypes*) in the supervised setting, i.e., the selection algorithm has access to the label information of the data points. Recently [30,22] have also explored the problem of prototype selection in the unsupervised setting, in which the selection algorithm has access only to the feature representation of the data points. They view the given dataset Y and a candidate prototype set P (subset of a source dataset X) as empirical distributions q and p , respectively. The prototype selection problem, therefore, is modeled as searching for a distribution p (corresponding to a set $P \subset X$ of data points, typically with a small cardinality) that is a good approximation of the distribution q . For example, [30,22] employ the maximum mean discrepancy (MMD) distance [21] to measure the similarity between the two distributions.

It is well-known that the MMD induces the “flat” geometry of reproducing kernel Hilbert space (RKHS) on the space of probability distributions, as it measures the distance between the mean embeddings of distributions in the RKHS of a universal kernel [47,21]). The individuality of data points is also lost while computing distance between mean embeddings in the MMD setting. The optimal transport (OT) framework, on the other hand, provides a natural metric for comparing probability distributions while respecting the underlying geometry of the data [51,39]. Over the last few years, OT distances (also known as the Wasserstein distances) have found widespread use in several machine learning applications such as image retrieval [44], shape interpolation [48], domain adaptation [6,7,28,37], supervised learning [18,27], and generative model training [1], among others. The *transport plan*, learned while computing the OT distance between the source and target distributions, is the joint distribution between the source and the target distributions. Compared to the MMD, the OT distances enjoy several advantages such as being faithful to the ground metric (geometry over the space of probability distributions) and identifying correspondences at the fine grained level of individual data points via the transport plan.

In this paper, we focus on the unsupervised prototype selection problem and view it from the perspective of the optimal transport theory. To this end, we propose a novel framework for **S**election of **P**rototypes using the **O**ptimal **T**ransport theory or the **SPOT** framework for searching a subset P from a source dataset X (i.e., $P \subset X$) that best represents a target set Y . We employ the Wasserstein distance to estimate the closeness between the distribution representing a candidate set P and set Y . Unlike the typical OT setting, the source distribution (representing P) is unknown in SPOT and needs to be learned along with the transport plan. The prototype selection problem is modeled as learning an empirical source distribution p (representing set X) that has the minimal Wasserstein distance with the empirical target distribution (representing set Y). Additionally, we constrain p to have a small support set (which represents $P \subset X$). The learned distribution p is also indicative of the relative importance of the prototypes in P in representing Y . Our main contributions are as follows.

- We propose a novel prototype selection framework, SPOT, based on the OT theory.
- We prove that the objective function of the proposed optimization problem in SPOT is submodular, which leads to a tight approximation guarantee of $(1 - e^{-1})$ using greedy approximation algorithms [38]. The computations in the proposed greedy algorithm can be implemented efficiently.

- We explain the popular k-medoids clustering [43] formulation as a special case of SPOT formulation (when the source and the target datasets are the same). We are not aware of any prior work that describes such a connection though the relation between Wasserstein distance minimization and k-means is known [4,11].
- Our empirical results show that the proposed algorithm outperforms existing baselines on several real-world datasets. The optimal transport framework allows our approach to seamlessly work in settings where the source (X) and the target (Y) datasets are from different domains.

The outline of the paper is as follows. We provide a brief review of the optimal transport setting, the prototype selection setting, and key definitions in the submodular optimization literature in Section 2. The proposed SPOT framework and algorithms are presented in Section 3. We discuss how SPOT relates to existing works in Section 4. The empirical results are presented in Section 5. We conclude the paper in Section 6. The proofs and additional results on datasets are available in our extended version [23].

2 Background

2.1 Optimal transport (OT)

Let $X := \{\mathbf{x}_i\}_{i=1}^m$ and $Y := \{\mathbf{y}_j\}_{j=1}^n$ be i.i.d. samples from the source and the target distributions p and q , respectively. In several applications, the true distributions are generally unknown. Their empirical estimates exist and can be employed as follows:

$$p := \sum_{i=1}^m \mathbf{p}_i \delta_{\mathbf{x}_i}, \quad q := \sum_{j=1}^n \mathbf{q}_j \delta_{\mathbf{y}_j}, \quad (1)$$

where the probability associated with samples \mathbf{x}_i and \mathbf{y}_j are \mathbf{p}_i and \mathbf{q}_j , respectively, and δ is the Dirac delta function. The vectors \mathbf{p} and \mathbf{q} lie on simplices Δ_m and Δ_n , respectively, where $\Delta_k := \{\mathbf{z} \in \mathbb{R}_+^k \mid \sum_i \mathbf{z}_i = 1\}$. The OT problem [29] aims at finding a transport plan γ (with the minimal transporting effort) as a solution to

$$\min_{\gamma \in \Gamma(\mathbf{p}, \mathbf{q})} \langle \mathbf{C}, \gamma \rangle, \quad (2)$$

where $\Gamma(\mathbf{p}, \mathbf{q}) := \{\gamma \in \mathbb{R}_+^{m \times n} \mid \gamma \mathbf{1} = \mathbf{p}; \gamma^\top \mathbf{1} = \mathbf{q}\}$ is the space of joint distribution between the source and the target marginals. Here, $\mathbf{C} \in \mathbb{R}_+^{m \times n}$ is the ground metric computed as $\mathbf{C}_{ij} = c(\mathbf{x}_i, \mathbf{y}_j)$ and the function $c : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+ : (\mathbf{x}, \mathbf{y}) \rightarrow c(\mathbf{x}, \mathbf{y})$ represents the *cost* of transporting a unit mass from source $\mathbf{x} \in \mathcal{X}$ to target $\mathbf{y} \in \mathcal{Y}$.

The optimization problem (2) is a linear program. Recently, [10] proposed an efficient solution for learning entropy regularized transport plan γ in (2) using the Sinkhorn algorithm [32]. For a recent survey on OT, please refer to [39].

2.2 Prototype selection

Selecting representative elements is often posed as identifying a subset P of size k from a set of items X (e.g., data points, features, etc.). The quality of selection is usually governed via a scoring function $f(P)$, which encodes the desirable properties of

prototypical samples. For instance, in order to obtain a compact yet informative subset P , the scoring function should discourage redundancy. Recent works [30,22] have posed prototype selection within the submodular optimization setting by maximizing a MMD based scoring function on the weights (\mathbf{w}) of the prototype elements:

$$l(\mathbf{w}) = \boldsymbol{\mu}^T \mathbf{w} - \frac{1}{2} \mathbf{w}^T \mathbf{K} \mathbf{w} \text{ s.t. } \|\mathbf{w}\|_0 \leq k. \quad (3)$$

Here, $\|\mathbf{w}\|_0$ is ℓ_0 norm of \mathbf{w} representing the number of non-zero values, the entries of the vector $\boldsymbol{\mu}$ contains the mean of the inner product for every source point with the target data points computed in the kernel embedding space, and \mathbf{K} is the Gram matrix of a universal kernel (e.g., Gaussian) corresponding to the source instances. The locations of non-zero values in \mathbf{w} , $\text{supp}(\mathbf{w}) = \{i : \mathbf{w}_i > 0\}$, known as its *support* correspond to the element indices that are chosen as prototypes, i.e. $P = \text{supp}(\mathbf{w})$. While the MMD-Critic method in [30] enforces that all non-zero entries in \mathbf{w} equal to $1/k$, the ProtoDash algorithm in [22] imposes non-negativity constraints and learns \mathbf{w} as part of the algorithm. Both propose greedy algorithms that effectively evaluate the *incremental* benefit of adding an element in the prototypical set P . In contrast to the MMD function in (3), to the best of our knowledge, ours is the first work which leverages the optimal transport (OT) framework to extract such compact representation. We prove that the proposed objective function is submodular, which ensures tight approximation guarantee using greedy approximate algorithms.

2.3 Submodularity

We briefly review the concept of submodular and weakly submodular functions, which we later use to prove key theoretical results.

Definition 1 (Submodularity and Monotonicity). Consider any two sets $A \subseteq B \subseteq [m]$. A set function $f(\cdot)$ is submodular if and only if for any $i \notin B$, $f(A \cup i) - f(A) \geq f(B \cup i) - f(B)$. The function is called monotone when $f(A) \leq f(B)$.

Submodularity implies diminishing returns where the incremental gain in adding a new element i to a set A is at least as high as adding to its superset B [19]. Another characterization of submodularity is via the submodularity ratio [16,12] defined as follows.

Definition 2 (Submodularity Ratio). Given two disjoint sets L and S , and a set function $f(\cdot)$, the submodularity ratio of $f(\cdot)$ for the ordered pair (L, S) is given by:

$$\alpha_{L,S} := \frac{\sum_{i \in S} [f(L \cup \{i\}) - f(L)]}{f(L \cup S) - f(L)}. \quad (4)$$

Submodularity ratio captures the increment in $f(\cdot)$ by adding the entire subset S to L , compared to summed gain of adding its elements individually to L . It is known that $f(\cdot)$ is submodular if and only if $\alpha_{L,S} \geq 1, \forall L, S$. In the case where $0 \leq \epsilon \leq \alpha_{L,S} < 1$ for an independent constant ϵ , $f(\cdot)$ is called *weakly submodular* [12].

We define submodularity ratio of a set P with respect to an integer s as follows:

$$\alpha_{P,s} := \max_{L,S:L \cap S = \emptyset, L \subseteq P, |S| \leq s} \alpha_{L,S}. \quad (5)$$

It should be emphasized that unlike the definition in [16, Equation 3], the above Equation (5) involves the \max operator instead of the \min . This specific form is later used to produce approximation bounds for the proposed approach (presented in Algorithm 1). Both (strongly) submodular and weakly submodular functions enjoy provable performance bounds when the set elements are selected incrementally and greedily [38,16,22].

3 SPOT framework

3.1 SPOT problem formulation

Let $X = \{\mathbf{x}_i\}_{i=1}^m$ be a set of m source points, $Y = \{\mathbf{y}_j\}_{j=1}^n$ be a target set of n data points, and $\mathbf{C} \in \mathbb{R}_+^{m \times n}$ represents the ground metric. Our aim is to select a small and weighted subset $P \subset X$ of size $k \ll m$ that best describes Y . To this end, we develop an optimal transport (OT) based framework for selection of prototypes. Traditionally, OT is defined as a minimization problem over the transport plan γ as in (2). In our setting, we pre-compute a *similarity* matrix $\mathbf{S} \in \mathbb{R}_+^{m \times n}$ from \mathbf{C} , for instance, as $\mathbf{S}_{ij} = \beta - \mathbf{C}_{ij}$ where $\beta > \|\mathbf{C}\|_\infty$. This allows to equivalently represent the OT problem (2) as a maximization problem with the objective function as $\langle \mathbf{S}, \gamma \rangle$. Treating it as a maximization problem enables to establish connection with submodularity and leverage standard greedy algorithms for its optimization [38].

We pose the problem of selecting a prototypical set as learning a sparse support empirical source distribution $w = \sum_{\mathbf{x}_i \in P} \mathbf{w}_i \delta_{\mathbf{x}_i}$ that has maximum closeness to the target distribution in terms of the optimal transport measure. Here, the weight $\mathbf{w} \in \Delta_m$, where $\Delta_m := \{\mathbf{z} \in \mathbb{R}_+^m \mid \sum_i \mathbf{z}_i = 1\}$. Consequently, \mathbf{w} denotes the relative importance of the samples. Hence, the constraint $|P| \leq k$ for the prototype set P translates to $|\text{supp}(\mathbf{w})| \leq k$ where $\text{supp}(\mathbf{w}) \subseteq P$. We evaluate the suitability of a candidate prototype set $P \subset X$ with an OT based measure on sets. To elaborate, index the elements in X from 1 to m and let $[m] := \{1, 2, \dots, m\}$ denote the first m natural numbers. Given any index set of prototypes $P \subseteq [m]$, define a set function $f : 2^{[m]} \rightarrow \mathbb{R}_+$ as:

$$f(P) := \max_{\mathbf{w}: \text{supp}(\mathbf{w}) \subseteq P} \max_{\gamma \in \Gamma(\mathbf{w}, \mathbf{q})} \langle \mathbf{S}, \gamma \rangle, \quad (6)$$

where $\mathbf{q} \in \Delta_n$ corresponds to the (given) weights of the target samples¹ in the empirical target distribution q as in (1). The learned transport plan γ in (6) is a joint distribution between the elements in P and Y , which may be useful in downstream applications requiring, e.g., barycentric mapping.

Our goal is to find that set P which maximizes $f(\cdot)$ subject to the cardinality constraint. To this end, the proposed SPOT problem is

$$P^* = \arg \max_{P \subseteq [m], |P| \leq k} f(P), \quad (7)$$

¹ In the absence of domain knowledge, uniform weights $\mathbf{q} = \mathbf{1}/n$ can be a default choice.

where $f(P)$ is defined in (6). The entries of the optimal weight vector \mathbf{w}^* corresponding to P^* in (7) indicate the importance of the prototypes in summarizing set Y . The SPOT (7) and the standard OT (2) settings are different as: (a) the source distribution w is learned as a part of the SPOT optimization problem formulation and (b) the source distribution w is enforced to have a sparse support of utmost size k so that the prototypes create a compact summary. In the next section, we analyze the objective function in the SPOT optimization problem (7), characterize it with a few desirable properties, and develop a computationally efficient greedy approximation algorithm.

3.2 Equivalent reduced representations of SPOT objective

Though the definition of $f(\cdot)$ in (6) involves maximization over two coupled variables \mathbf{w} and γ , it can be reduced to an equivalent optimization problem involving only γ (by eliminating \mathbf{w} altogether). To this end, let $k = |P|$ and denote \mathbf{S}_P a $k \times n$ sub-matrix of \mathbf{S} containing only those rows indexed by P . We then have the following lemma:

Lemma 3. *The set function $f(\cdot)$ in (6) can be equivalently defined as an optimization problem only over the transport plan, i.e.,*

$$f(P) = \max_{\gamma \in \Gamma_P(\mathbf{q})} \langle \mathbf{S}_P, \gamma \rangle, \quad (8)$$

where $\Gamma_P(\mathbf{q}) := \{\gamma \in \mathbb{R}_+^{k \times n} \mid \gamma^T \mathbf{1} = \mathbf{q}\}$. Let γ^* be an optimal solution of (8). Then, (\mathbf{w}^*, γ^*) is an optimal solution of (6) where $\mathbf{w}^* = \gamma^* \mathbf{1}$.

A closer look into the set function in (8) reveals that the optimization for γ can be done in parallel over the n target points, and its solution assumes a closed-form expression. It is worth noting that the constraint $\gamma^T \mathbf{1} = \mathbf{q}$ as well as the objective $\langle \mathbf{S}_P, \gamma \rangle$ decouple over each column of γ . Hence, (8) can be solved across the columns of variable γ independently, thereby allowing parallelism over the target set. In other words,

$$f(P) = \sum_{j=1}^n \max_{\gamma^j \in \mathbb{R}_+^k} \langle \mathbf{S}_P^j, \gamma^j \rangle, \text{ s.t. } \mathbf{1}^T \gamma^j = \mathbf{q}_j \forall j, \quad (9)$$

where \mathbf{S}_P^j and γ^j denote the j^{th} column vectors of the matrices \mathbf{S}_P and γ , respectively. Furthermore, if i_j denotes the location of the maximum value in the vector \mathbf{S}_P^j , then an optimal solution γ^* can be easily seen to inherit an extremely sparse structure with exactly one non-zero element in each column j at the row location i_j , i.e., $\gamma_{i_j, j}^* = \mathbf{q}_j, \forall j$ and 0 everywhere. So (9) can be reduced to

$$f(P) = \sum_{j=1}^n \mathbf{q}_j \max_{i \in P} \mathbf{S}_{ij}. \quad (10)$$

The above observation makes the computation $f(P)$ in (10) particularly suited when using GPUs. Further, due to this specific solution structure in (10), determining the function value for any incremental set is an inexpensive operation as shown below.

Lemma 4 (Fast incremental computation). *Given any set P and its function value $f(P)$, the value at the incremental selection $f(P \cup S)$ obtained by adding $s = |S|$ new elements to P , can be computed in $O(sn)$.*

Remark 5. By setting $P = \emptyset$ and $f(\emptyset) = 0$, $f(S)$ for any set S can be determined efficiently as discussed in Lemma 4.

3.3 SPOT optimization algorithms

As obtaining the global optimum subset P^* for the problem (7) is NP complete, we now present two approximation algorithms for SPOT: SPOTsimple and SPOTgreedy.

SPOTsimple: a fast heuristic algorithm SPOTsimple is an extremely fast heuristic that works as follows. For every source point \mathbf{x}_i , SPOTsimple determines the indices of target points $\mathcal{T}_i = \{j : \mathbf{S}_{ij} \geq \mathbf{S}_{\tilde{i}j} \text{ for all } \tilde{i} \neq i\}$ that have the highest similarity to \mathbf{x}_i compared to other source points. In other words, it solves (10) with $P = [m]$, i.e., no cardinality constraint, to determine the initial transport plan γ where $\gamma_{ij} = \mathbf{q}_j$ if $j \in \mathcal{T}_i$ and 0 everywhere else. It then computes the source weights as $\mathbf{w} = \gamma \mathbf{1}$ with each entry $w_i = \sum_{j \in \mathcal{T}_i} \mathbf{q}_j$. The top- k source points based on the weights \mathbf{w} are chosen as the prototype set P . The final transport plan γ_P is recomputed using (10) over P . The total computational cost incurred by SPOTsimple for selecting k prototypes is $O(mn)$.

SPOTgreedy: a greedy and incremental prototype selection algorithm As we discuss later in our experiments (section 5), though SPOTsimple is computationally very efficient, its accuracy of prototype selection is sensitive to the skewness of class instances in the target distribution. When the samples from different classes are uniformly represented in the target set, SPOTsimple is indeed able to select prototypes from the source set that are representative of the target. However, when the target is skewed and the class distributions are no longer uniform, SPOTsimple primarily chooses from the dominant class leading to biased selection and poor performance (see Figure 2(a)).

To this end, we present our method of choice SPOTgreedy, detailed in Algorithm 1, that leverages the following desirable properties of the function $f(\cdot)$ in (10) to *greedily and incrementally* build the prototype set P . For choosing k prototypes, SPOTgreedy costs $O(mnk/s)$. As most operations in SPOTgreedy involve basic matrix manipulations, the practical implementation cost of SPOTgreedy is considerably low.

Lemma 6 (Submodularity). *The set function $f(\cdot)$ defined in (10) is monotone and submodular [36].*

The submodularity of $f(\cdot)$ enables to provide provable approximation bounds for greedy element selections in SPOTgreedy. The algorithm begins by setting the current selection $P = \emptyset$. Without loss of generality, we assume $f(\emptyset) = 0$ as $f(\cdot)$ is monotonic. In each iteration, it determines those s elements from the remainder set $[m] \setminus P$, denoted by S , that when individually added to P result in maximum incremental gain. This can be implemented efficiently as discussed in Lemma 4. Here $s \geq 1$ is the user parameter

Algorithm 1 SPOTgreedy

Input: sparsity level k or lower bound ϵ on increment in $f(\cdot)$, X , Y , s , and \mathbf{q} .
Initialize $P = \emptyset$
while $|P| \leq k$ or increment in objective $\geq \epsilon$. **do**
 Define vector β with entries $\beta_i = f(P \cup \{i\}) - f(P)$, $\forall i \in [m] \setminus P$.
 $S =$ Set of indices of top s largest elements in β .
 $P = P \cup S$.
end while
 $\gamma_P = \arg \max_{\gamma \in \Gamma_P(\mathbf{q})} \langle \mathbf{S}_P, \gamma \rangle$; $\mathbf{w}_P = \gamma_P \mathbf{1}$.
Return $P, \gamma_P, \mathbf{w}_P$.

that decides the number of elements chosen in each iteration. The set S is then added to P . The algorithm proceeds for $\lceil \frac{k}{s} \rceil$ iterations to select k prototypes. As function $f(\cdot)$ in (8) is both monotone and submodular, it has the characteristic of diminishing returns. Hence, an alternative stopping criterion could be the minimum expected increment ϵ in the function value at each iteration. The algorithm stops when the increment in the function value is below the specified threshold ϵ .

Approximation guarantee for SPOTgreedy. We note the following result on the upper bound on the submodularity ratio (4). Let $s = |S|$. When $f(\cdot)$ is monotone, then

$$\alpha_{L,S} \leq \frac{\sum_{i \in S} [f(L \cup \{i\}) - f(L)]}{\max_{i \in S} [f(L \cup \{i\}) - f(L)]} \leq s \quad (11)$$

and hence $\alpha_{P,s} \leq s$. In particular, $s = 1$ implies $\alpha_{P,1} = 1$, as for any $L \subseteq P$, $\alpha_{L,S} = 1$ when $|S| = 1$. Our next result provides the performance bound for the proposed SPOTgreedy algorithm.

Theorem 7 (Performance bounds for SPOTgreedy). *Let P be the final set returned by the SPOTgreedy method described in Algorithm 1. Let $\alpha = \alpha_{P,s}$ be the submodularity ratio of the set P w.r.t. s . If P^* is the optimal set of k elements that maximizes $f(\cdot)$ in the SPOT optimization problem (7), then*

$$f(P) \geq f(P^*) \left[1 - e^{-\frac{1}{\alpha}}\right] \geq f(P^*) \left[1 - e^{-\frac{1}{s}}\right]. \quad (12)$$

When $s = 1$ we recover the known approximation guarantee of $(1 - e^{-1})$ [38].

3.4 k-medoids as a special case of SPOT

Consider the specific setting where the source and the target datasets are the same, i.e., $X = Y$. Let $n = |X|$ and $\mathbf{q}_j = 1/n$ having uniform weights on the samples. Selecting a prototypical set $P \subset X$ is in fact a data summarization problem of choosing few representative exemplars from a given set of n data points, and can be thought as an output of a clustering method where P contains the the cluster centers. A popular clustering method is the k-medoids algorithm that ensures the cluster centers are exemplars

chosen from actual data points [43]. As shown in [36], the objective function for the k-medoids problem is

$$g(P) = \frac{1}{n} \sum_{j=1}^n \max_{\mathbf{z} \in P} l(\mathbf{z}, \mathbf{x}_j),$$

where $l(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{S}_{ij}$ defines the similarity between the respective data points. Comparing it against (10) gives a surprising connection that the *k-medoids algorithm is a special case of learning an optimal transport plan with a sparse support in the setting where the source and target distributions are the same*. Though the relation between OT and k-means is discussed in [4,11], we are not cognizant of any prior works that explains k-medoids from the lens of optimal transport. However, the notion of *transport* loses its relevance as there is no distinct target distribution to which the source points need to be transported. It should be emphasized that the connection with k-medoids is only in the limited case where the source and target distributions are the same. Hence, the popular algorithms that solve the k-medoids problem [46] like PAM, CLARA, and CLARANS cannot be applied in the general setting when the distributions are *different*.

4 Related works and discussion

As discussed earlier, recent works [30,22] view the unsupervised prototype selection problem as searching for a set $P \subset X$ whose underlying distribution is similar to the one corresponding to the target dataset Y . However, instead of the true source and target distributions, only samples from them are available. In such a setting, φ -divergences [9] e.g., the total variation distance and KL-divergence, among others require density estimation or space-partitioning/bias-correction techniques [47], which can be computationally prohibitive in higher dimensions. Moreover, they may be agnostic to the natural geometry of the ground metric. The maximum mean discrepancy (MMD) metric (3) employed by [30,22], on the other hand, can be computed efficiently but does not faithfully lift the ground metric of the samples [17].

We propose an optimal transport (OT) based prototype selection approach. OT framework respects the intrinsic geometry of the space of the distributions. Moreover, there is an additional flexibility in the choice of the ground metric, e.g., ℓ_1 -norm distance, which need not be a (universal) kernel induced function sans which the distribution approximation guarantees of MMD may no longer be applicable [21]. Solving the classical OT problem (2) is known to be computationally more expensive than computing MMD. However, our setting differs from the classical OT setup, as the source distribution is also learned in (6). As shown in Lemmas 3 & 4, the joint learning of the source distribution and the optimal transport plan has an equivalent but computationally efficient reformulation (8).

Using OT is also favorable from a theoretical standpoint. Though the MMD function in [30] is proven to be submodular, it is only under restricted conditions like the choice of kernel matrix and equal weighting of prototypes. The work in [22] extends [30] by allowing for unequal weights and eliminating any additional conditions on the kernel, but forgoes submodularity as the resultant MMD objective (3) is only weakly submodular. In this backdrop, the SPOT objective function (7) is submodular without requiring

any further assumptions. It is worth noting that submodularity leads to a tighter approximation guarantee of $(1 - e^{-1})$ using greedy approximation algorithms [38], whereas the best greedy based approximation for weak submodular functions (submodularity ratio of $\alpha < 1$) is only $(1 - e^{-\alpha})$ [16]. A better theoretical approximation of the OT based subset selection encourages the selection of better quality prototypes.

5 Experiments

We evaluate the generalization performance and computational efficiency of our algorithms against state-of-the-art on several real-world datasets. The codes are available at <https://pratikjawanpuria.com>. The following methods are compared.

- **MMD-Critic** [30]: it uses a maximum mean discrepancy (MMD) based scoring function. All the samples are weighted equally in the scoring function.
- **ProtoDash** [22]: it uses a *weighted* MMD based scoring function. The learned weights indicate the importance of the samples.
- **SPOTsimple**: our fast heuristic algorithm described in Section 3.3.
- **SPOTgreedy**: our greedy and incremental algorithm (Algorithm 1).

Following [2,30,22], we validate of the quality of the representative samples selected by different prototype selection algorithms via the performance of the corresponding *nearest prototype classifier*. Let X and Y represent source and target datasets containing different class distributions and let $P \subseteq X$ be a candidate representative set of the target Y . The quality of P is evaluated by classifying the target set instances with 1-nearest neighbour (1-NN) classifier parameterized by the elements in P . The class information of the samples in P is made available during this evaluation stage. Such classifiers can achieve better generalization performance than the standard 1-NN classifier due to reduction of noise overfitting [8] and have been found useful for large scale classification problems [54,50].

5.1 Prototype selection within same domain

We consider the following benchmark datasets.

- **ImageNet** [45]: we use the popular subset corresponding to ILSVRC 2012-2017 competition. The images have 2048 dimensional deep features [24].
- **MNIST** [34] is a handwritten digit dataset consisting of greyscale images of digits $\{0, \dots, 9\}$. The images are of 28×28 pixels.
- **USPS** dataset [25] consists of handwritten greyscale images of $\{0, \dots, 9\}$ digits represented as 16×16 pixels.
- **Letter** dataset [15] consists of images of twenty-six capital letters of the English alphabets. Each letter is represented as a 16 dimensional feature vector.
- **Flickr** [49] is the Yahoo/Flickr Creative Commons multi-label dataset consisting of descriptive tags of various real-world outdoor/indoor images.

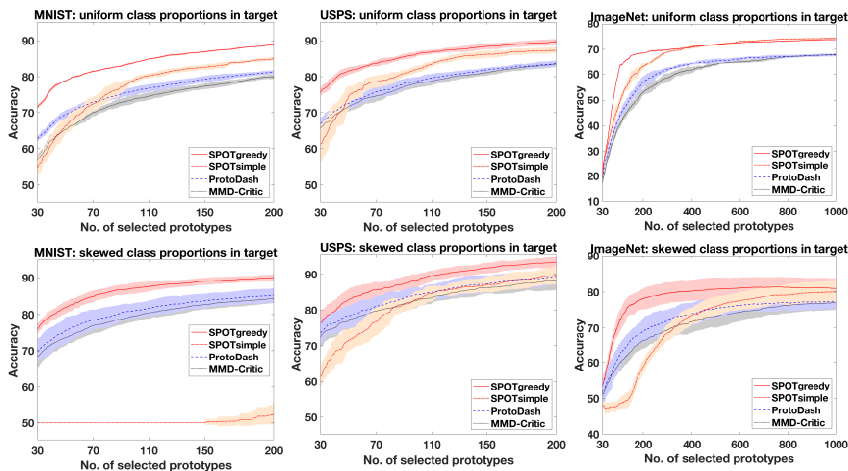


Fig. 1. Performance of different prototype selection algorithms. The standard deviation for every k is represented as a lighter shaded band around the mean curve corresponding to each method. [Top row] all the classes have uniform representation in the target set. [Bottom row] the challenging skewed setting where a randomly chosen class represents 50% of the target set (while the remaining classes together uniformly represent the rest).

Results on the Letter and Flickr datasets are discussed in the extended version [23].

Experimental setup. In the first set of experiments, all the classes are equally represented in the target set. In second set of experiments, the target sets are skewed towards a randomly chosen class, whose instances (digit/letter) form $z\%$ of the target set and the instances from the other classes uniformly constitute the remaining $(100 - z)\%$. For a given dataset, the source set is same for all the experiments and uniformly represents all the classes. Results are averaged over ten randomized runs. More details on the experimental set up are given in the extended version [23].

Results. Figure 1 (top row) shows the results of the first set of experiments on MNIST, USPS, and ImageNet. We plot the test set accuracy for a range of top- k prototypes selected. We observe that the proposed SPOTgreedy outperforms ProtoDash and MMD-Critic over the whole range of k . Figure 1 (bottom row) shows the results when samples of a (randomly chosen) class constitutes 50% of the target set. SPOTgreedy again dominates in this challenging setting. We observe that in several instances, SPOTgreedy opens up a significant performance gap even with only a few selected prototypes. The average running time on CPU of algorithms on the ImageNet dataset are: 55.0s (SPOTgreedy), 0.06s (SPOTsimple), 911.4s (ProtoDash), and 710.5s (MMD-Critic). We observe that both our algorithms, SPOTgreedy and SPOTsimple, are much faster than both ProtoDash and MMD-Critic.

Figure 2(a) shows that SPOTgreedy achieves the best performance on different skewed versions of the MNIST dataset (with $k = 200$). Interestingly, in cases where the target distribution is either uniform or heavily skewed, our heuristic non-incremental

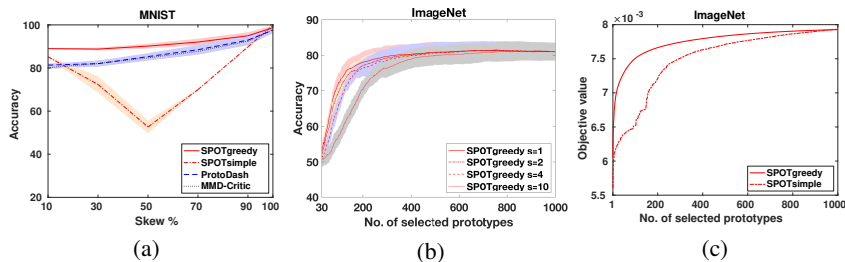


Fig. 2. (a) Comparisons of different algorithms in representing targets with varying skew percentage of a MNIST digit; (b) Performance of our SPOTgreedy algorithm with varying subset selection size s on the ImageNet dataset; (c) Comparison of the objective value (7) obtained by the proposed algorithms SPOTgreedy and SPOTsimple for various values of k . SPOTgreedy consistently obtains a better approximation.

algorithm SPOTsimple can select prototypes that match the target distribution well. However, in the harder setting when skewness of class instances in the target dataset varies from 20% to 80%, SPOTsimple predominantly selects the skewed class leading to a poor performance.

In Figure 2(b), we plot the performance of SPOTgreedy for different choices of s (which specifies the number of elements chosen simultaneously in each iteration). We consider the setting where the target has 50% skew of one of the ImageNet digits. Increasing s proportionally decreases the computational time as the number of iterations $\lceil \frac{k}{s} \rceil$ steadily decreases with s . However, choosing few elements simultaneously generally leads to better target representation. We note that between $s = 1$ and $s = 10$, the degradation in quality is only marginal even when we choose as few as 110 prototypes and the performance gap continuously narrows with more prototype selection. However, the time taken by SPOTgreedy with $s = 10$ is 5.7s, which is almost the expected 10x speedup compared to SPOTgreedy with $s = 1$ which takes 55.0s. In this setting, we also compare the qualitative performance of the proposed algorithms in solving Problem (7). Figure 2(c) shows the objective value obtained after every selected prototype on ImageNet. SPOTgreedy consistently obtains a better objective than SPOTsimple, showing the benefit of the greedy and incremental selection approach.

Identifying criticisms for MNIST. We further make use of the prototypes selected by SPOTgreedy to identify *criticisms*. These are data points belonging to the region of input space not well explained by prototypes and are farthest away from them. We use a witness function similar to [30, Section 3.2]. The columns of Figure 3(b) visualizes the few chosen criticisms, one for each of the 10 datasets containing samples of the respective MNIST digits. It is evident that the selected data points are indeed outliers for the corresponding digit class. Since the criticisms are those points that are maximally dissimilar from the prototypes, it is also a reflection on how well the prototypes of SPOTgreedy represent the underlying class as seen in Figure 3(a), where in each column we plot the selected prototypes for a dataset comprising one of the ten digits.

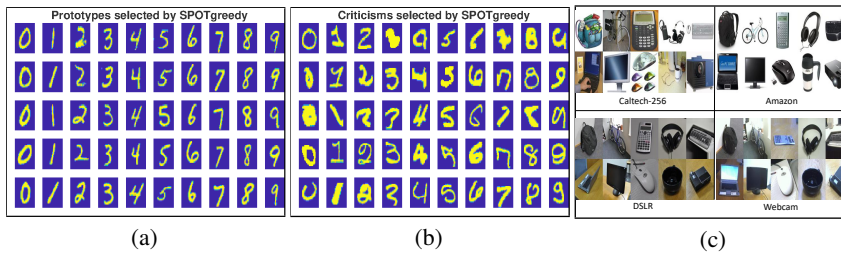


Fig. 3. (a) Prototypes selected by SPOTgreedy for the dataset containing one of the ten MNIST digits (column-wise); (b) Criticisms chosen by SPOTgreedy for the dataset containing one of the ten MNIST digits (column-wise); (c) Example images representing the ten classes in the four domains of the Office-Caltech dataset [20].

Table 1. Accuracy obtained on the Office-Caltech dataset.

Task	MMD-Critic	MMD-Critic+OT	ProtoDash	ProtoDash+OT	SPOTsimple	SPOTgreedy
$A \rightarrow C$	73.98	78.16	70.23	72.28	82.62	83.60
$A \rightarrow D$	75.16	72.61	77.71	71.97	80.25	82.80
$A \rightarrow W$	51.53	62.71	48.81	58.64	62.37	75.59
$C \rightarrow A$	83.71	86.17	83.82	87.25	71.92	90.03
$C \rightarrow D$	70.06	75.16	71.34	70.70	75.80	89.17
$C \rightarrow W$	49.83	54.92	46.44	53.56	70.85	82.03
$D \rightarrow A$	82.85	85.21	83.39	83.82	91.00	90.89
$D \rightarrow C$	78.25	78.34	75.40	79.41	85.38	86.27
$D \rightarrow W$	80.00	84.41	85.08	86.10	75.59	92.20
$W \rightarrow A$	71.60	78.56	68.38	74.71	87.03	84.99
$W \rightarrow C$	67.20	75.76	65.86	74.60	74.06	83.12
$W \rightarrow D$	92.36	96.18	88.54	89.81	86.62	94.90
Average	73.04	77.35	72.08	75.24	78.36	86.30

5.2 Prototype selection from different domains

Section 5.1 focused on settings where the source and the target datasets had similar/dissimilar class distributions. We next consider a setting where the source and target datasets additionally differ in feature distribution, e.g., due to covariate shift [41].

Figure 2(c) shows examples from the classes of the Office-Caltech dataset [20], which has images from four domains: Amazon (online website), Caltech (image dataset), DSLR (images from a DSLR camera), and Webcam (images from a webcam). Images from the same class vary across the four domains due to several factors such as different background, lighting conditions, etc. The number of data points in each domain is: 958 (A: Amazon), 1123 (C: Caltech), 157 (D: DSLR), and 295 (W: Webcam). The number of instances per class per domain ranges from 8 to 151. DeCAF6 features [14] of size 4096 are used for all the images. We design the experiment similar to Section 5.1 by considering each domain, in turn, as the source or the target. There are twelve different tasks where task $A \rightarrow W$ implies that Amazon and Webcam are the source and the target domains, respectively. The total number of selected prototypes is 20.

Results. Table 1 reports the accuracy obtained on every task. We observe that our SPOTgreedy significantly outperforms MMD-Critic and ProtoDash. This is because SPOTgreedy learns both the prototypes as well as the transport plan between the prototypes and the target set. The transport plan allows the prototypes to be transported to the target domain via the barycentric mapping, a characteristic of the optimal transport framework. SPOTgreedy is also much better than SPOTsimple due to its superior incremental nature of prototype selection. We also empower the non-OT based baselines for the domain adaptation setting as follows. After selecting the prototypes via a baseline, we learn an OT plan between the selected prototypes and the target data points by solving the OT problem (2). The distribution of the prototypes is taken to be the normalized weights obtained by the baseline. This ensures that the prototypes selected by MMD-Critic+OT, and ProtoDash+OT are also transported to the target domain. Though we observe marked improvements in the performance of MMD-Critic+OT and ProtoDash+OT, the proposed SPOTgreedy and SPOTsimple still outperform them.

6 Conclusion

We have looked at the prototype selection problem from the viewpoint of optimal transport. In particular, we show that the problem is equivalent to learning a sparse source distribution w , whose probability values w_i specify the relevance of the corresponding prototype in representing the given target set. After establishing connections with submodularity, we proposed the SPOTgreedy algorithm that employs incremental greedy selection of prototypes and comes with (i) deterministic theoretical guarantees, (ii) simple implementation with updates that are amenable to parallelization, and (iii) excellent performance on different benchmarks.

Future works: A few interesting generalizations and research directions are as follows.

- Our k -prototype selection problem (7) may be viewed as learning a ℓ_0 -norm regularized (fixed-support) Wasserstein barycenter of a single distribution. Extending it to learning sparse Wasserstein barycenter of multiple distributions may be useful in applications like model compression, noise removal, etc.
- With the Gromov-Wasserstein (GW) distance [35,40], the OT distance has been extended to settings where the source and the target distributions do not share the same feature and metric space. Extending SPOT with the GW-distances is useful when the source and the target domains share similar concepts/categories/classes but are defined over different feature spaces.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: ICML (2017)
2. Bien, J., Tibshirani, R.: Prototype Selection for Interpretable Classification. *Annals of Applied Statistics* **5**(4), 2403–2424 (2011)
3. Bien, J., Tibshirani, R.: Hierarchical clustering with prototypes via minimax linkage. *Journal of the American Statistical Association* **106**(495), 1075–1084 (2011)
4. Canas, G., Rosasco, L.: Learning probability measures with respect to optimal transport metrics. In: *NeurIPS* (2012)

5. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., Elhadad, N.: Intelligible models for healthcare. In: SIGKDD (2015)
6. Courty, N., Flamary, R., Habrard, A., Rakotomamonjy, A.: Joint distribution optimal transportation for domain adaptation. In: NeurIPS (2017)
7. Courty, N., Flamary, R., Tuia, D., Rakotomamonjy, A.: Optimal transport for domain adaptation. TPAMI **39**(9), 1853–1865 (2017)
8. Crammer, K., Gilad-Bachrach, R., Navot, A., Tishby, N.: Margin analysis of the LVQ algorithm. In: NeurIPS (2002)
9. Csiszár, I.: A class of measures of informativity of observation channels. Periodica Mathematica Hungarica **2**(1), 191–213 (1972)
10. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. In: NeurIPS (2013)
11. Cuturi, M., Doucet, A.: Fast computation of Wasserstein barycenters. In: ICML (2014)
12. Das, A., Kempe, D.: Submodular meets Spectral: Greedy Algorithms for Subset Selection, Sparse Approximation and Dictionary Selection. In: ICML (2011)
13. Dhurandhar, A., Gurumoorthy, K.S.: Classifier invariant approach to learn from positive-unlabeled data. In: IEEE ICDM (2020)
14. Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: DeCAF: A deep convolutional activation feature for generic visual recognition. In: ICML (2014)
15. Dua, D., Graff, C.: UCI machine learning repository (2017)
16. Elenberg, E., Khanna, R., Dimakis, A.G., Negahban, S.: Restricted Strong Convexity Implies Weak Submodularity. Annals of Statistics **46**, 3539–3568 (2018)
17. Feydy, J., Séjourné, T., Vialard, F.X., Amari, S., Trounev, A., Peyré, G.: Interpolating between optimal transport and MMD using Sinkhorn divergences. In: AISTATS (2018)
18. Frogner, C., Zhang, C., Mobahi, H., Araya-Polo, M., Poggio, T.: Learning with a wasserstein loss. In: NeurIPS (2015)
19. Fujishige, S.: Submodular functions and optimization. Elsevier (2005)
20. Gong, B., Shi, Y., Sha, F., Grauman, K.: Geodesic flow kernel for unsupervised domain adaptation. In: CVPR (2012)
21. Gretton, A., Borgwardt, K.M., Rasch, M., Schölkopf, B., Smola, A.J.: A kernel two-sample test. Journal of Machine Learning Research **13**(25), 723–773 (2012)
22. Gurumoorthy, K.S., Dhurandhar, A., Cecchi, G., Aggarwal, C.: Efficient data representation by selecting prototypes with importance weights. In: IEEE ICDM (2019)
23. Gurumoorthy, K.S., Jawanpuria, P., Mishra, B.: SPOT: A framework for selection of prototypes using optimal transport. Tech. rep., arXiv preprint arXiv:2103.10159 (2021)
24. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
25. Hull, J.: A database for handwritten text recognition research. TPAMI **16**(5), 550–554 (1994)
26. Idé, T., Dhurandhar, A.: Supervised Item Response Models for Informative Prediction. Knowl. Inf. Syst. **51**(1), 235–257 (2017)
27. Jawanpuria, P., Dev, S., Mishra, B.: Efficient robust optimal transport: Formulations and algorithms. Tech. rep., arXiv preprint arXiv:2010.11852 (2020)
28. Jawanpuria, P., Meghwanshi, M., Mishra, B.: Geometry-aware domain adaptation for unsupervised alignment of word embeddings. In: ACL (2020)
29. Kantorovich, L.: On the translocation of masses. Doklady of the Academy of Sciences of the USSR **37**, 199–201 (1942)
30. Kim, B., Khanna, R., Koyejo, O.: Examples are not Enough, Learn to Criticize! Criticism for Interpretability. In: NeurIPS (2016)
31. Kim, B., Rudin, C., Shah, J.: The Bayesian case model: A generative approach for case-based reasoning and prototype classification. In: NeurIPS (2014)

32. Knight, P.A.: The Sinkhorn-Knopp algorithm: Convergence and applications. *SIAM Journal on Matrix Analysis and Applications* **30**(1), 261–275 (2008)
33. Koh, P.W., Liang, P.: Understanding black-box predictions via influence functions. In: *ICML* (2017)
34. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (November 1998)
35. Mémoli, F.: Gromov-Wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics* **11**(4), 417–487 (2011)
36. Mirzasoleiman, B., Karbasi, A., Sarkar, R., Krause, A.: Distributed submodular maximization. *Journal of Machine Learning Research* **17**(235), 1–44 (2016)
37. Nath, J.S., Jawanpuria, P.: Statistical optimal transport posed as learning kernel mean embedding. In: *NeurIPS* (2020)
38. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An Analysis of Approximations for Maximizing Submodular Set Functions. *Mathematical Programming* **14**, 265–294 (1978)
39. Peyré, G., Cuturi, M.: Computational optimal transport. *Foundations and Trends in Machine Learning* **11**(5-6), 355–607 (2019)
40. Peyré, G., Cuturi, M., Solomon, J.: Gromov-Wasserstein averaging of kernel and distance matrices. In: *ICML* (2016)
41. Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.: *Dataset Shift in Machine Learning*. The MIT Press (2009)
42. Ribeiro, M., Singh, S., Guestrin, C.: "Why Should I Trust You?" Explaining the Predictions of Any Classifier. In: *SIGKDD* (2016)
43. Rousseeuw, P.J., Kaufman, L.: *Finding Groups in Data: An Introduction to Cluster Analysis*. John Wiley & Sons, Inc. (2009)
44. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover's distance as a metric for image retrieval. *IJCV* **40**(2), 99–121 (2000)
45. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *IJCV* **115**(3), 211–252 (2015)
46. Schubert, E., Rousseeuw, P.J.: Faster k-Medoids clustering: Improving the PAM, CLARA, and CLARANS algorithms. In: *International Conference on Similarity Search and Applications* (2019)
47. Smola, A., Gretton, A., Song, L., Schölkopf, B.: A Hilbert space embedding for distributions. In: *International Conference on Algorithmic Learning Theory* (2007)
48. Solomon, J., De Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., Guibas, L.: Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics* **34**(4), 66:1–66:11 (2015)
49. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: Yfcc100m: The new data in multimedia research. *Communications of ACM* **59**(2), 64–73 (2016)
50. Tibshirani, R., Hastie, T., Narasimhan, B., Chu, G.: Diagnosis of multiple cancer types by shrunken centroids of gene expression. *PNAS* **99**(10), 6567–6572 (2002)
51. Villani, C.: *Optimal Transport: Old and New*. Springer Verlag (2009)
52. Wei, K., Iyer, R., Bilmes, J.: Submodularity in data subset selection and active learning. In: *ICML* (2015)
53. Weiser, M.: Programmers use slices when debugging. *Comm. ACM* **25**(7), 446–452 (1982)
54. Wohlhart, P., Köstinger, M., Donoser, M., Roth, P., Bischof, H.: Optimizing 1-nearest prototype classifiers. In: *CVPR* (2013)
55. Yeh, C.K., Kim, J., Yen, I.E.H., Ravikumar, P.K.: Representer point selection for explaining deep neural networks. In: *NeurIPS* (2018)