# Subset Selection by Pareto Optimization with Recombination 

Chao Qian, ${ }^{1 *}$ Chao Bian, ${ }^{1}$ Chao Feng ${ }^{2}$<br>${ }^{1}$ National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China<br>${ }^{2}$ School of Computer Science and Technology, University of Science and Technology of China, Hefei 230027, China<br>qianc@lamda.nju.edu.cn, chaobian12@gmail.com, chaofeng@mail.ustc.edu.cn


#### Abstract

Subset selection, i.e., to select a limited number of items optimizing some given objective function, is a fundamental problem with various applications such as unsupervised feature selection and sparse regression. By employing a multi-objective evolutionary algorithm (EA) with mutation only to optimize the given objective function and minimize the number of selected items simultaneously, the recently proposed POSS algorithm achieves state-of-the-art performance for subset selection. In this paper, we propose the PORSS algorithm by incorporating recombination, a characterizing feature of EAs, into POSS. We prove that PORSS can achieve the optimal polynomial-time approximation guarantee as POSS when the objective function is monotone, and can find an optimal solution efficiently in some cases whereas POSS cannot. Extensive experiments on unsupervised feature selection and sparse regression show the superiority of PORSS over POSS. Our analysis also theoretically discloses that recombination from diverse solutions can be more likely than mutation alone to generate various variations, thereby leading to better exploration; this may be of independent interest for understanding the influence of recombination.


## Introduction

This paper considers a general problem, i.e., subset selection, which is to select a subset of size at most $k$ from a total set of $n$ items for maximizing (or minimizing) some given objective function $f$. This problem arises in various real-world applications, such as maximum coverage (Feige 1998), sparse regression (Miller 2002), influence maximization (Kempe, Kleinberg, and Tardos 2003), sensor placement (Krause, Singh, and Guestrin 2008), document summarization (Lin and Bilmes 2011) and unsupervised feature selection (Farahat, Ghodsi, and Kamel 2011), to name a few.

Subset selection is generally NP-hard, and much efforts have been devoted to developing polynomial-time approximation algorithms. The greedy algorithm, which iteratively selects one item with the largest marginal gain, has been shown to be a good approximation solver. When the involved objective function $f$ satisfies the monotone property, the greedy algorithm can achieve the $\left(1-e^{-\gamma}\right)$ -

[^0]approximation guarantee, where $\gamma$ is the submodularity ratio measuring how close $f$ is to submodularity (Das and Kempe 2011). Particularly, for submodular objective functions, $\gamma=1$ and the approximation guarantee becomes $1-1 / e$, which is optimal, i.e., cannot be improved by any polynomial-time algorithm (Nemhauser and Wolsey 1978). Harshaw et al. (2019) have recently proved that the general approximation guarantee of $\left(1-e^{-\gamma}\right)$ is also optimal.

Based on Pareto optimization, Qian et al. (2015) proposed the POSS algorithm for subset selection. The idea is to reformulate subset selection as a bi-objective optimization problem maximizing the given objective and minimizing the subset size simultaneously, then solve the problem by a multiobjective EA (MOEA), and finally select the best solution with size at most $k$ from the generated solution set. It has been shown that POSS can achieve the optimal polynomialtime approximation guarantee, $1-e^{-\gamma}$, and can be significantly better than the greedy algorithm in applications, e.g., unsupervised feature selection and sparse regression. Moreover, POSS is robust against uncertainties (Qian et al. 2017; Roostapour et al. 2019), and easily distributed for large-scale tasks (Qian et al. 2016; 2018; Qian 2019).

The optimization engine of POSS is the employed MOEA, which iteratively reproduces new solutions for solving the reformulated bi-objective problem. For EAs, mutation and recombination (or called crossover) are two popular operators for reproduction (Bäck 1996); the former changes one solution randomly whereas the latter mixes up two or more solutions. POSS applies mutation only and has performed well, while recombination, as a core feature of EAs, may be helpful to further improve its performance.

In this paper, we propose the PORSS algorithm for subset selection by introducing recombination into POSS. Two common recombination operators are considered: one-point recombination and uniform recombination. In theory, we prove that for subset selection with monotone objective functions, PORSS can achieve the optimal polynomial-time approximation guarantee, $1-e^{-\gamma}$; for one concrete example of subset selection, PORSS can be significantly faster than POSS to find an optimal solution. We also conduct experiments on the applications of unsupervised feature selection and sparse regression with various real-world data sets, showing that within the same running time, PORSS can almost always achieve better performance than POSS.

Note that recombination is understood only at preliminary level, though there are great efforts devoted to analyzing its influence, e.g., (Neumann and Theile 2010; Doerr et al. 2013; Qian, Yu, and Zhou 2013; Oliveto and Witt 2014; Sudholt 2017; Dang et al. 2018). Our analysis theoretically discloses that recombining diverse solutions is more likely than mutation to generate various variations, and thus to escape from local optima; this may help to understand this kind of operator.

## Subset Selection

Given a ground set $V=\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}$, we study the functions $f: 2^{V} \rightarrow \mathbb{R}$ over subsets of $V$. A set function $f$ is monotone if $\forall S \subseteq T, f(S) \leq f(T)$. Assume w.l.o.g. that monotone functions are normalized, i.e., $f(\emptyset)=0$. A set function $f$ is submodular (Nemhauser, Wolsey, and Fisher 1978) if $\forall S \subseteq T \subseteq V$,

$$
f(T)-f(S) \leq \sum_{v \in T \backslash S}(f(S \cup\{v\})-f(S))
$$

For a general set function $f$, the notion of submodularity ratio in Definition 1 is used to measure to what extent $f$ has the submodular property. When $f$ is monotone, it holds that (1) $\forall S, l: 0 \leq \gamma_{S, l}(f) \leq 1$, and (2) $f$ is submodular iff $\forall S, l: \gamma_{S, l}(f)=1$.
Definition 1 (Submodularity Ratio (Das and Kempe 2011)). The submodularity ratio of a set function $f: 2^{V} \rightarrow \mathbb{R}$ with respect to a set $S \subseteq V$ and a parameter $l \geq 1$ is
$\gamma_{S, l}(f)=\min _{L \subseteq S, T:|T| \leq l, T \cap L=\emptyset} \frac{\sum_{v \in T}(f(L \cup\{v\})-f(L))}{f(L \cup T)-f(L)}$.
The subset selection problem as presented in Definition 2 is to select a subset $S$ of $V$ such that a given objective $f$ is maximized with the constraint $|S| \leq k$. For a monotone function $f$, the greedy algorithm, which iteratively adds one item with the largest marginal gain until $k$ items are selected, can achieve an approximation guarantee of ( $1-e^{-\gamma_{S, k}(f)}$ ) (Das and Kempe 2011), where $S$ is the subset output by the greedy algorithm. The optimality of this approximation guarantee was known only in the case where $\gamma_{S, k}(f)=1$, i.e., $f$ is submodular (Nemhauser and Wolsey 1978), and has recently been proved in the general case (Harshaw et al. 2019).
Definition 2 (Subset Selection). Given all items $V=$ $\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}$, an objective function $f$ and a budget $k$, to find a subset of at most $k$ items maximizing $f$, i.e.,

$$
\begin{equation*}
\arg \max _{S \subseteq V} f(S) \quad \text { s.t. } \quad|S| \leq k \tag{1}
\end{equation*}
$$

Here are two applications of subset selection with monotone, but not necessarily submodular, objective functions, that will be studied in this paper. Unsupervised feature selection as presented in Definition 3 is to select at most $k$ columns from a matrix $\mathbf{A}$ to best approximate $\mathbf{A}$. Some notations: $(\cdot)^{+}$: Moore-Penrose inverse of a matrix; $\|\cdot\|_{F}$ : Frobenius norm of a matrix; $|\cdot|$ : number of columns of a matrix. The goodness of approximation is measured by the sum of squared errors between the original matrix $\mathbf{A}$ and the
approximation $\mathbf{S S}^{+} \mathbf{A}$, where $\mathbf{S S}^{+}$is the projection matrix onto the space spanned by the columns of $\mathbf{S}$. Note that a submatrix of $\mathbf{A}$ can be seen as a subset of all columns of $\mathbf{A}$.
Definition 3 (Unsupervised Feature Selection). Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a budget $k$, to find a submatrix $\mathbf{S}$ of $\mathbf{A}$ with at most $k$ columns minimizing $\left\|\mathbf{A}-\mathbf{S S}^{+} \mathbf{A}\right\|_{F}^{2}$, i.e.,

$$
\arg \min _{\mathbf{S}:} \text { a submatrix of } \mathbf{A}\left\|\mathbf{A}-\mathbf{S S}^{+} \mathbf{A}\right\|_{F}^{2} \quad \text { s.t. } \quad|\mathbf{S}| \leq k
$$

For the ease of theoretical treatment, this minimization problem is often equivalently reformulated as a maximization problem (Bhaskara et al. 2016; Ordozgoiti, Canaval, and Mozo 2018):

$$
\arg \max _{\mathbf{S}: \text { a submatrix of } \mathbf{A}}\left\|\mathbf{S S}^{+} \mathbf{A}\right\|_{F}^{2} \quad \text { s.t. } \quad|\mathbf{S}| \leq k .
$$

Sparse regression (Miller 2002) as presented in Definition 4 is to find a sparse approximation solution to the linear regression problem. Note that $S$ and its index set $\left\{i \mid v_{i} \in\right.$ $S\}$ are not distinguished for notational convenience, and all variables are assumed w.l.o.g. to be normalized to have expectation 0 and variance 1 .
Definition 4 (Sparse Regression). Given observation variables $V=\left\{v_{1}, \ldots, v_{n}\right\}$, a predictor variable $z$ and a budget $k$, to find at most $k$ variables from $V$ maximizing the squared multiple correlation (Johnson and Wichern 2007), i.e.,

$$
\arg \max _{S \subseteq V} R_{z, S}^{2}=1-\mathrm{MSE}_{z, S} \quad \text { s.t. } \quad|S| \leq k
$$

where $\mathrm{MSE}_{z, S}$ denotes the mean squared error, i.e.,

$$
\operatorname{MSE}_{z, S}=\min _{\boldsymbol{\alpha} \in \mathbb{R}|S|} \mathbb{E}\left[\left(z-\sum_{i \in S} \alpha_{i} v_{i}\right)^{2}\right]
$$

## The POSS Algorithm

Based on Pareto optimization, a new algorithm POSS for subset selection has been proposed (Friedrich and Neumann 2015; Qian, Yu, and Zhou 2015). Note that a subset $S$ of $V$ can be represented by a binary vector $\boldsymbol{x} \in\{0,1\}^{n}$, where $x_{i}=1$ iff the item $v_{i} \in S$, and we will not distinguish them for notational convenience. POSS reformulates the original problem Eq. (1) as a bi-objective minimization problem:

$$
\begin{equation*}
\arg \min _{\boldsymbol{x} \in\{0,1\}^{n}} \quad\left(f_{1}(\boldsymbol{x}), f_{2}(\boldsymbol{x})\right) \tag{2}
\end{equation*}
$$

where

$$
f_{1}(\boldsymbol{x})=\left\{\begin{array}{ll}
+\infty, & |\boldsymbol{x}| \geq 2 k \\
-f(\boldsymbol{x}), & \text { otherwise }
\end{array}, \quad f_{2}(\boldsymbol{x})=|\boldsymbol{x}|\right.
$$

Thus, POSS maximizes the original objective function $f$ and minimizes the subset size $|\boldsymbol{x}|$ simultaneously. By setting $f_{1}$ to $+\infty$ for $|x| \geq 2 k$, overly infeasible solutions, i.e., solutions with large constraint violation, are excluded.

To compare solutions in bi-objective optimization, POSS uses the domination relationship. For two solutions $\boldsymbol{x}$ and $\boldsymbol{x}^{\prime}, \boldsymbol{x}$ weakly dominates $\boldsymbol{x}^{\prime}$, denoted as $\boldsymbol{x} \preceq \boldsymbol{x}^{\prime}$, if $f_{1}(\boldsymbol{x}) \leq$ $f_{1}\left(\boldsymbol{x}^{\prime}\right) \wedge f_{2}(\boldsymbol{x}) \leq f_{2}\left(\boldsymbol{x}^{\prime}\right) ; \boldsymbol{x}$ dominates $\boldsymbol{x}^{\prime}$, denoted as $\boldsymbol{x} \prec$ $\boldsymbol{x}^{\prime}$, if $\boldsymbol{x} \preceq \boldsymbol{x}^{\prime}$ and either $f_{1}(\boldsymbol{x})<f_{1}\left(\boldsymbol{x}^{\prime}\right)$ or $f_{2}(\boldsymbol{x})<f_{2}\left(\boldsymbol{x}^{\prime}\right)$; they are incomparable, if neither $\boldsymbol{x} \preceq \boldsymbol{x}^{\prime}$ nor $\boldsymbol{x}^{\prime} \preceq \boldsymbol{x}$.

POSS employs a simple MOEA with mutation only, which is slightly modified from the GSEMO algorithm (Giel

```
Algorithm 1 POSS Algorithm
Input: \(V=\left\{v_{1}, \ldots, v_{n}\right\}\); objective \(f: 2^{V} \rightarrow \mathbb{R}\); budget \(k\)
Parameter: the number \(T\) of iterations
Output: a subset of \(V\) with at most \(k\) items
Process:
    Let \(\boldsymbol{x}=0^{n}, P=\{\boldsymbol{x}\}\) and \(t=0\);
    while \(t<T\) do
        Select \(\boldsymbol{x}\) from \(P\) randomly;
        Apply bit-wise mutation on \(\boldsymbol{x}\) to generate \(\boldsymbol{x}^{\prime}\);
        if \(\nexists \boldsymbol{z} \in P\) such that \(\boldsymbol{z} \prec \boldsymbol{x}^{\prime}\) then
            \(P=\left(P \backslash\left\{\boldsymbol{z} \in P \mid \boldsymbol{x}^{\prime} \preceq \boldsymbol{z}\right\}\right) \cup\left\{\boldsymbol{x}^{\prime}\right\}\)
        end if
        \(t=t+1\)
    end while
    return \(\arg \max _{\boldsymbol{x} \in P,|\boldsymbol{x}| \leq k} f(\boldsymbol{x})\)
```

2003; Laumanns, Thiele, and Zitzler 2004), to solve the biobjective problem Eq. (2). As described in Algorithm 1, it starts from $0^{n}$ representing the empty set, and iteratively tries to improve the solutions in the population $P$ (lines 2-9). In each iteration, a solution $\boldsymbol{x}$ is selected from $P$ uniformly at random, and used to generate a new solution $\boldsymbol{x}^{\prime}$ by the bit-wise mutation operator, presented as follows:

Bit-wise mutation: flip each bit of a solution $\boldsymbol{x} \in\{0,1\}^{n}$ independently with probability $1 / n$.
The newly generated solution $\boldsymbol{x}^{\prime}$ is then used to update $P$ in lines 5-7, making $P$ contain only non-dominated solutions generated-so-far. That is, if $x^{\prime}$ is not dominated by any solution in $P$ (line 5), it will be added into $P$, and meanwhile those archived solutions weakly dominated by $\boldsymbol{x}^{\prime}$ will be deleted (line 6). After running $T$ iterations, the best solution w.r.t. the original problem Eq. (1) is selected from $P$ in line 10 as the final output solution.

For subset selection with monotone objective functions, POSS has been proved to achieve the same general approximation guarantee as the greedy algorithm in polynomial expected running time, i.e., to achieve the optimal polynomialtime approximation guarantee (Qian, Yu, and Zhou 2015). Furthermore, it has been empirically shown that POSS can achieve significantly better performance than the greedy algorithm in some applications, e.g., unsupervised feature selection (Feng, Qian, and Tang 2019) and sparse regression (Qian, Yu, and Zhou 2015).

## The PORSS Algorithm

To reproduce new solutions in each iteration, POSS applies the mutation operator, which simulates the mutation phenomena in DNA transformation. It is known that recombination is another popular operator for reproduction, which simulates the chromosome exchange phenomena in zoogamy reproduction, and typically appears in various real EAs, e.g., the popularly used algorithm NSGA-II (Deb et al. 2002).

In this section, we propose a new Pareto Optimization algorithm with Recombination for Subset Selection, briefly called PORSS. As described in Algorithm 2, PORSS employs recombination and mutation together, rather than mutation only, to generate new solutions in each iteration. In

```
Algorithm 2 PORSS Algorithm
Input: \(V=\left\{v_{1}, \ldots, v_{n}\right\}\); objective \(f: 2^{V} \rightarrow \mathbb{R}\); budget \(k\)
Parameter: the number \(T\) of iterations
Output: a subset of \(V\) with at most \(k\) items
Process:
    Let \(\boldsymbol{x}=0^{n}, P=\{\boldsymbol{x}\}\) and \(t=0\);
    while \(t<T\) do
        Select \(\boldsymbol{x}, \boldsymbol{y}\) from \(P\) randomly with replacement;
        Apply recombination on \(\boldsymbol{x}, \boldsymbol{y}\) to generate \(\boldsymbol{x}^{\prime}, \boldsymbol{y}^{\prime}\);
        Apply bit-wise mutation on \(\boldsymbol{x}^{\prime}, \boldsymbol{y}^{\prime}\) to generate \(\boldsymbol{x}^{\prime \prime}, \boldsymbol{y}^{\prime \prime}\);
        for each \(\boldsymbol{q} \in\left\{\boldsymbol{x}^{\prime \prime}, \boldsymbol{y}^{\prime \prime}\right\}\)
            if \(\nexists \boldsymbol{z} \in P\) such that \(\boldsymbol{z} \prec \boldsymbol{q}\) then
                \(P=(P \backslash\{\boldsymbol{z} \in P \mid \boldsymbol{q} \preceq \boldsymbol{z}\}) \cup\{\boldsymbol{q}\}\)
            end if
        end for
        \(t=t+1\)
    end while
    return \(\arg \max _{\boldsymbol{x} \in P,|\boldsymbol{x}| \leq k} f(\boldsymbol{x})\)
```

line 3 , two solutions are selected randomly from the population $P$ with replacement, and then recombined to generate new solutions in line 4 . We consider two commonly used recombination operators:

One-point recombination: select $i \in\{1,2, \ldots, n\}$ randomly, and exchange the first $i$ bits of two solutions;

Uniform recombination: exchange each bit of two solutions independently with probability $1 / 2$.
For example, for solutions $0^{n}$ and $1^{n}$, two new solutions $0^{n / 2} 1^{n / 2}$ and $1^{n / 2} 0^{n / 2}$ can be generated by one-point recombination with probability $1 / n$, and by uniform recombination with probability $\left(1 / 2^{n}\right) \cdot 2$, where the factor 2 is included due to the symmetry. In line 5, the two solutions generated by recombination are further mutated to generate another two ones, which are used to update the population $P$ in lines 6-10.

## Influence of Recombination

To understand the influence of recombination intuitively, we compare the distribution of the number of bits flipped with/without recombination. Suppose two solutions $\boldsymbol{x}, \boldsymbol{y}$ selected in line 3 of Algorithm 2 have Hamming distance $d$, denoted as $H(\boldsymbol{x}, \boldsymbol{y})=d$. Let $\boldsymbol{x}^{\prime \prime}, \boldsymbol{y}^{\prime \prime}$ denote the two solutions generated by recombination and mutation in line 5 . We analyze the probability for $\boldsymbol{x}^{\prime \prime}$ or $\boldsymbol{y}^{\prime \prime}$ to have Hamming distance $j$ with $\boldsymbol{x}$, denoted as $Q(d, j)$. That is,
$Q(d, j)=\mathrm{P}\left(H\left(\boldsymbol{x}, \boldsymbol{x}^{\prime \prime}\right)=j \vee H\left(\boldsymbol{x}, \boldsymbol{y}^{\prime \prime}\right)=j \mid H(\boldsymbol{x}, \boldsymbol{y})=d\right)$.
For one-point and uniform recombination, we use the notations $Q_{o}(d, j)$ and $Q_{u}(d, j)$, respectively. Let $\boldsymbol{z}_{m}$ denote the solution generated from a solution $\boldsymbol{z}$ by bit-wise mutation. By turning off recombination, i.e., deleting line 4 of Algorithm 2 , we analyze the corresponding probability
$Q_{m}(d, j)=\mathrm{P}\left(H\left(\boldsymbol{x}, \boldsymbol{x}_{m}\right)=j \vee H\left(\boldsymbol{x}, \boldsymbol{y}_{m}\right)=j \mid H(\boldsymbol{x}, \boldsymbol{y})=d\right)$.
In the following, we compare $Q_{o}(d, j), Q_{u}(d, j)$ with $Q_{m}(d, j)$, to examine the influence of recombination.

Given a solution $\boldsymbol{z}$ with Hamming distance $i$ from $\boldsymbol{x}$, let $q_{i, j}$ denote the probability for the Hamming distance to become $j$ by bit-wise mutation, i.e.,

$$
q_{i, j}=\mathrm{P}\left(H\left(\boldsymbol{x}, \boldsymbol{z}_{m}\right)=j \mid H(\boldsymbol{x}, \boldsymbol{z})=i\right)
$$

For $Q_{m}(d, j)$, as $H(\boldsymbol{x}, \boldsymbol{x})=0$ and $H(\boldsymbol{x}, \boldsymbol{y})=d$, we have

$$
Q_{m}(d, j)=q_{0, j}+q_{d, j}-q_{0, j} \cdot q_{d, j} .
$$

By uniform recombination, $\boldsymbol{x}, \boldsymbol{y}$ exchange $i$ different bits with probability $\binom{d}{i}(1 / 2)^{i}(1 / 2)^{d-i}=\binom{d}{i}(1 / 2)^{d}$, generating two solutions $\boldsymbol{x}^{\prime}, \boldsymbol{y}^{\prime}$ where $H\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)=i$ and $H\left(\boldsymbol{x}, \boldsymbol{y}^{\prime}\right)=$ $d-i$. Note that $\boldsymbol{x}$ and $\boldsymbol{y}$ have totally $d$ different bits. Considering the mutation behavior on $\boldsymbol{x}^{\prime}, \boldsymbol{y}^{\prime}$, we have

$$
Q_{u}(d, j)=\sum_{i=0}^{d}\binom{d}{i} \frac{1}{2^{d}} \cdot\left(q_{i, j}+q_{d-i, j}-q_{i, j} \cdot q_{d-i, j}\right)
$$

Consider one-point recombination. $\forall 1 \leq i \leq d$, there exists $l \geq i$ such that exchanging the first $l$ bits of $\boldsymbol{x}, \boldsymbol{y}$ can generate two solutions $\boldsymbol{x}^{\prime}, \boldsymbol{y}^{\prime}$ where $H\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)=i$ and $H\left(\boldsymbol{x}, \boldsymbol{y}^{\prime}\right)=$ $d-i$. Thus, we have

$$
Q_{o}(d, j) \geq \frac{1}{n} \sum_{i=1}^{d}\left(q_{i, j}+q_{d-i, j}-q_{i, j} \cdot q_{d-i, j}\right)
$$

Because it is sufficient to keep all bits unchanged in mutation, $q_{j, j} \geq(1-1 / n)^{n} \geq 1 /(2 e)$. Thus, we have
(a) $\forall j \leq d: Q_{o}(d, j) \geq 1 / n \cdot q(j, j)=\Omega(1 / n)$;
(b) $\forall j \leq d: Q_{u}(d, j) \geq\binom{ d}{j} \frac{1}{2^{d}} \cdot q(j, j)=\Omega\left(\binom{d}{j} / 2^{d}\right)$.

By analyzing $q_{0, j}$ and $q_{d, j}$, we can derive that, $\exists 0<j_{0} \leq d$,
(c.1) for $j<j_{0}: \Omega\left((1 / j)^{j}\right) \leq Q_{m}(d, j) \leq O\left((e / j)^{j}\right)$; (c.2) for $j \geq j_{0}$ :

$$
\Omega\left(\frac{1}{d-j} \cdot \frac{d}{n}\right)^{d-j} \leq Q_{m}(d, j) \leq O\left(\frac{e}{d-j} \cdot \frac{d}{n}\right)^{d-j}
$$

The detailed analysis for $Q_{m}(d, j)$ is provided in the supplementary material due to space limitations.

According to (c.1) and (c.2), the number of bits flipped by mutation only is strongly concentrated around two extreme values, 0 and $d$. When $j$ increases from 0 to $j_{0}$ or decreases from $d$ to $j_{0}, Q_{m}(d, j)$ decays super-exponentially. Particularly, for $j=d / 2, q(0, d / 2) \leq\binom{ n}{d / 2}(1 / n)^{d / 2} \leq 1 /(d / 2)$ !, $q(d, d / 2) \leq\binom{ d}{d / 2}(1 / n)^{d / 2} \leq 1 /(d / 2)!$, and thus,

$$
\begin{equation*}
Q_{m}(d, d / 2) \leq \frac{2}{(d / 2)!} \leq \frac{2}{e(d /(2 e))^{d / 2}} \leq\left(\frac{2 e}{d}\right)^{d / 2} \tag{3}
\end{equation*}
$$

where the second inequality holds by Stirling's formula. According to (b), the number of bits flipped by uniform recombination and mutation is concentrated around $d / 2$, but $Q_{u}(d, j)$ is always lower bounded by $\Omega\left(1 / 2^{d}\right)$, which is much greater than $Q_{m}(d, d / 2)$ in Eq. (3) when $d$ is large. According to (a), $\forall j \leq d: Q_{o}(d, j) \geq \Omega(1 / n)$, implying that the number of bits flipped by one-point recombination and mutation is relatively uniformly distributed.

Therefore, from diverse solutions, i.e., when $d$ is large, recombination can ease flipping any number of bits, and may lead to better exploration and thus a better ability of escaping from local optima. The advantage of recombination will be verified by theoretical analysis and empirical study.

## Theoretical Analysis

As introduced before, the greedy algorithm and POSS can achieve the optimal polynomial-time approximation guarantee for subset selection with monotone objective functions. A natural question is whether PORSS can keep the optimal approximation. We give the positive answer by proving Theorem 1, i.e., PORSS achieves the approximation guarantee of $\left(1-e^{-\gamma_{\text {min }}}\right)$ in polynomial expected running time. Let OPT denote the optimal function value. The proof is inspired by the analysis of POSS (Qian, Yu, and Zhou 2015).
Lemma 1. (Qian et al. 2016) Let $f:\{0,1\}^{n} \rightarrow \mathbb{R}^{+}$be a monotone function. For any $\boldsymbol{x} \in\{0,1\}^{n}$, there exists one item $v \notin \boldsymbol{x}$ such that

$$
f(\boldsymbol{x} \cup\{v\})-f(\boldsymbol{x}) \geq \frac{\gamma_{\boldsymbol{x}, k}}{k}(\mathrm{OPT}-f(\boldsymbol{x}))
$$

Theorem 1. For subset selection with any monotone $f$, the expected number of iterations until PORSS with one-point or uniform recombination finds a solution $\boldsymbol{x}$ with $|\boldsymbol{x}| \leq k$ and $f(\boldsymbol{x}) \geq\left(1-e^{-\gamma_{\min }}\right) \cdot$ OPT is polynomial, where $\gamma_{\min }=$ $\min _{\boldsymbol{x}:|\boldsymbol{x}|=k-1} \gamma_{\boldsymbol{x}, k}$.
Proof. Let $J_{\max }$ be the maximum value of $j \in\{0,1, \ldots, k\}$ such that in the population $P$, there exists a solution $\boldsymbol{x}$ with $|\boldsymbol{x}| \leq j$ and $f(\boldsymbol{x}) \geq\left(1-\left(1-\gamma_{\min } / k\right)^{j}\right) \cdot$ OPT. That is,

$$
\begin{aligned}
J_{\max }= & \max \{j \in\{0,1, \ldots, k\} \mid \exists \boldsymbol{x} \in P: \\
& \left.|\boldsymbol{x}| \leq j \wedge f(\boldsymbol{x}) \geq\left(1-\left(1-\gamma_{\min } / k\right)^{j}\right) \cdot \mathrm{OPT}\right\} .
\end{aligned}
$$

We only need to analyze the expected number of iterations until $J_{\max }=k$, which implies that there exists one solution $\boldsymbol{x} \in P$ satisfying that $|\boldsymbol{x}| \leq k$ and $f(\boldsymbol{x}) \geq(1-(1-$ $\left.\left.\gamma_{\min } / k\right)^{k}\right) \cdot$ OPT $\geq\left(1-e^{-\gamma_{\min }}\right) \cdot$ OPT.

As PORSS starts from $0^{n}, J_{\max }$ is initially 0 . Assume that currently $J_{\max }=i<k$. Let $\boldsymbol{x}$ denote a solution corresponding to $J_{\max }=i$, i.e., $|\boldsymbol{x}| \leq i$ and $f(\boldsymbol{x}) \geq$ $\left(1-\left(1-\gamma_{\min } / k\right)^{i}\right) \cdot$ OPT. First, $J_{\max }$ will not decrease. This is because deleting $\boldsymbol{x}$ from $P$ in line 8 of Algorithm 2 implies that $\boldsymbol{x}$ is weakly dominated by the newly included solution $\boldsymbol{q}$, satisfying that $|\boldsymbol{q}| \leq|\boldsymbol{x}| \leq i$ and $f(\boldsymbol{q}) \geq f(\boldsymbol{x}) \geq$ $\left(1-\left(1-\gamma_{\min } / k\right)^{i}\right) \cdot$ OPT.

Next, we analyze the probability of increasing $J_{\max }$ in one iteration. Consider the case that the two selected solutions in line 3 of Algorithm 2 are both $x$, occurring with probability $(1 /|P|) \cdot(1 /|P|)$ due to uniform selection with replacement. For two identical solutions, either one-point or uniform recombination in line 4 makes no changes. Thus, in line $5, x$ is used to generate a new solution by bit-wise mutation, and this process is implemented twice independently. For bit-wise mutation on $\boldsymbol{x}$, according to Lemma 1 , a new solution $\boldsymbol{x}^{\prime}$ satisfying $f\left(\boldsymbol{x}^{\prime}\right)-f(\boldsymbol{x}) \geq\left(\gamma_{\boldsymbol{x}, k} / k\right)$. (OPT $-f(\boldsymbol{x})$ ) can be generated by flipping only one specific 0 bit of $\boldsymbol{x}$ (i.e., adding a specific item into $\boldsymbol{x}$ ), occurring with probability $(1 / n)(1-1 / n)^{n-1} \geq 1 /(e n)$. As

$$
\begin{aligned}
f(\boldsymbol{x}) \geq(1 & \left.-\left(1-\gamma_{\min } / k\right)^{i}\right) \cdot \text { OPT, we have } \\
f\left(\boldsymbol{x}^{\prime}\right) & \geq\left(1-\gamma_{\boldsymbol{x}, k} / k\right) \cdot f(\boldsymbol{x})+\left(\gamma_{\boldsymbol{x}, k} / k\right) \cdot \mathrm{OPT} \\
& \geq\left(1-\left(1-\gamma_{\boldsymbol{x}, k} / k\right)\left(1-\gamma_{\min } / k\right)^{i}\right) \cdot \mathrm{OPT} \\
& \geq\left(1-\left(1-\gamma_{\min } / k\right)^{i+1}\right) \cdot \mathrm{OPT} .
\end{aligned}
$$

Note that the last inequality holds by $\gamma_{\boldsymbol{x}, k} \geq \gamma_{\text {min }}$, because $|\boldsymbol{x}|<k$ and $\gamma_{\boldsymbol{x}, k}$ decreases with $\boldsymbol{x}$. As $\boldsymbol{x}$ is mutated twice independently in line 5 , such a new solution $x^{\prime}$ can be generated with probability at least $1-(1-1 /(e n))^{2}=$ $2 /(e n)-1 /(e n)^{2}$. It is clear that $\left|\boldsymbol{x}^{\prime}\right|=|\boldsymbol{x}|+1 \leq i+1$. Then, $\boldsymbol{x}^{\prime}$ will be added into $P$; otherwise, $\boldsymbol{x}^{\prime}$ must be dominated by one archived solution in line 7 of Algorithm 2, and this implies that $J_{\text {max }}$ has been larger than $i$, contradicting with the assumption $J_{\max }=i$. After adding $x^{\prime}$ into $P, J_{\max } \geq i+1$. Thus, $J_{\max }$ can increase by at least 1 in one iteration with probability at least $\left(1 /|P|^{2}\right) \cdot\left(2 /(e n)-1 /(e n)^{2}\right)$. By the procedure of updating the population $P$ in lines 6-10, the solutions in $P$ must be incomparable. Thus, each value of one objective can correspond to at most one solution in $P$. Because the solutions with $|x| \geq 2 k$ have $+\infty$ value on the first objective, they must be excluded from $P$, and thus, $|\boldsymbol{x}| \in$ $\{0,1, \ldots, 2 k-1\}$, implying $|P| \leq 2 k$. We can now conclude that the probability of increasing $J_{\max }$ in one iteration is at least $\left(1 /\left(4 k^{2}\right)\right) \cdot\left(2 /(e n)-1 /(e n)^{2}\right)=\Omega\left(1 /\left(k^{2} n\right)\right)$.

The above analysis shows that $J_{\max }$ will not decrease, but can increase with probability $\Omega\left(1 /\left(k^{2} n\right)\right)$ in one iteration. Thus, the expected number of iterations until $J_{\text {max }}$ increases by at least 1 is $O\left(k^{2} n\right)$. For $J_{\max }=k$, it requires to increase $J_{\text {max }}$ by at most $k$ times, implying that the expected number of iterations until finding a solution with the desired approximation guarantee is $O\left(k^{3} n\right)$, which is polynomial.

Next, by an illustrative example of subset selection, we prove that PORSS can perform much better than the greedy algorithm and POSS. As presented in Definition 5, the best subset of size $(i+1)$ can be generated from the best subset of size $i$ by adding one specific item, and the only exception is the best subset of size $k$, i.e., the optimal solution, which differs greatly from the best subsets of other sizes. This example represents subset selection problems where decisions have to be made in sequence to some extent.

## Definition 5. The objective function $f$ satisfies that

(1) $\forall 0 \leq i \leq n-1: f\left(\boldsymbol{x}^{i}\right)<f\left(\boldsymbol{x}^{i+1}\right)$;
(2) if $|\boldsymbol{x}|=i \neq k$, then $\forall \boldsymbol{x} \neq \boldsymbol{x}^{i}: f(\boldsymbol{x})<f\left(\boldsymbol{x}^{i}\right)$;
(3) if $|\boldsymbol{x}|=k$, then $\forall \boldsymbol{x} \notin\left\{\boldsymbol{x}^{*}, \boldsymbol{x}^{k}\right\}: f(\boldsymbol{x})<f\left(\boldsymbol{x}^{k}\right)<f\left(\boldsymbol{x}^{*}\right)$, where $\boldsymbol{x}^{i}=1^{i} 0^{n-i}, \boldsymbol{x}^{*}=0^{k} 1^{k} 0^{n-2 k}$ and $k \leq n / 2$.

It is clear that the optimal solution is $\boldsymbol{x}^{*}=0^{k} 1^{k} 0^{n-2 k}$, and each $\boldsymbol{x}^{i}=1^{i} 0^{n-i}$ is the best solution for size $i$ except that $\boldsymbol{x}^{k}=1^{k} 0^{n-k}$ is the runner-up for size $k$. Due to the greedy nature, the greedy algorithm finds $\boldsymbol{x}^{0}, \boldsymbol{x}^{1}, \ldots, \boldsymbol{x}^{k}$ sequentially, implying that $\boldsymbol{x}^{*}$ cannot be found.
Lemma 2. For subset selection with $f$ in Definition 5, the greedy algorithm cannot find the optimal solution.

Lemmas 3 to 5 show the expected number of iterations of POSS and PORSS until finding the optimal solution. The
detailed proofs are provided in the supplementary material due to space limitations, and we introduce the proof intuition here. Both POSS and PORSS can find the solutions $\left\{\boldsymbol{x}^{0}, \boldsymbol{x}^{1}, \ldots, \boldsymbol{x}^{2 k-1}\right\}$ efficiently. After that, the population will always consist of these solutions before finding the optimal one. Because the Hamming distance between $\boldsymbol{x}^{i}$ and $\boldsymbol{x}^{*}$ is at least $k$, the probability of generating the optimal solution $\boldsymbol{x}^{*}$ by mutation is at most $(1 / n)^{k}$, and thus, POSS is inefficient. For PORSS, recombination from the two diverse solutions $\boldsymbol{x}^{0}=0^{n}$ and $\boldsymbol{x}^{2 k-1}=1^{2 k-1} 0^{n-2 k+1}$ can generate the solution $0^{k} 1^{k-1} 0^{n-2 k+1}$ through exchanging their first $k$ bits, occurring with a large probability, i.e., $1 / n$ or $\left(1 / 2^{2 k-1}\right) \cdot 2$, by one-point or uniform recombination; the subsequent mutation operator can generate $x^{*}$ by flipping only one specific bit, i.e., the $(2 k)$-th bit, occurring with probability $(1 / n)(1-1 / n)^{n-1}$. Thus, PORSS can be efficient. Note that the reason for the effectiveness of recombination is consistent with that found in the last section.
Lemma 3. For subset selection with $f$ in Definition 5, the expected number of iterations until POSS finds the optimal solution is at least $\left(n /\left(3 e^{4 e k}\right)\right)^{k}$.
Lemma 4. For subset selection with $f$ in Definition 5, the expected number of iterations until PORSS with one-point recombination finds the optimal solution is at most $6 e k^{2} n^{2}$.
Lemma 5. For subset selection with $f$ in Definition 5, the expected number of iterations until PORSS with uniform recombination finds the optimal solution is at most $5 k^{2} 4^{k} n$.

Let $\mathbb{E}\left[T_{m}\right], \mathbb{E}\left[T_{o}\right]$ and $\mathbb{E}\left[T_{u}\right]$ denote the expected number of iterations of POSS, PORSS with one-point and uniform recombination, respectively, for finding the optimal solution. Considering that the budget $k$ is usually not large in real applications, we make the following observations.
Remark 1. According to Lemmas 3 to 5, we have
(1) if $k \leq O(1)$, then $\mathbb{E}\left[T_{m}\right] \geq \Omega\left(n^{k}\right)$,

$$
\mathbb{E}\left[T_{o}\right] \leq O\left(n^{2}\right) \text { and } \mathbb{E}\left[T_{u}\right] \leq O(n)
$$

(2) if $\omega(1) \leq k \leq(\log n) / 20$, then $\mathbb{E}\left[T_{m}\right] \geq n^{k / 4}$,

$$
\mathbb{E}\left[T_{o}\right] \leq n^{2}(\log n)^{2} \text { and } \mathbb{E}\left[T_{u}\right] \leq n^{1.1}(\log n)^{2}
$$

In other words, when $k$ is a constant, i.e., $O(1)$, PORSS is polynomially faster than POSS, and the gap increases with $k$; when $k$ continues to increase to $\omega(1)$, the gap becomes super-polynomially large.

## Empirical Study

In this section, we empirically compare PORSS, POSS and the greedy algorithm on the applications of unsupervised feature selection and sparse regression with various realworld data sets. ${ }^{1}$ PORSS using one-point and uniform recombination are denoted by PORSS $_{o}$ and PORSS $_{u}$, respectively. Note that some common algorithms, e.g., IterFS (Ordozgoiti, Canaval, and Mozo 2018) for unsupervised feature selection and lasso (Tibshirani 1996) for sparse regression,

[^1]Table 1: Unsupervised feature selection: the error ratio (the smaller, the better) of the compared algorithms on ten data sets for $k=8$. The mean $\pm$ std. is reported for randomized algorithms. In each data set, the smallest values are bolded. The count of direct win denotes the number of data sets on which POSS has a smaller error ratio than the corresponding algorithm ( 1 tie is counted as 0.5 win ), where significant cells by the sign-test with confidence level 0.05 are bolded.

| Data set | (\#inst, \#feat) | OPT | Greedy | POSS | PORSS $_{o}$ | PORSS $_{u}$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| Sonar | $(208,60)$ | 1.353 | 1.429 | $1.371 \pm 0.007$ | $\mathbf{1 . 3 5 8} \pm \mathbf{0 . 0 0 6}$ | $1.363 \pm 0.010$ |
| phishing | $(11055,68)$ | 1.166 | 1.223 | $1.168 \pm 0.006$ | $\mathbf{1 . 1 6 6} \pm \mathbf{0 . 0 0 0}$ | $1.167 \pm 0.003$ |
| Hill-Valley | $(606,100)$ | - | 1.544 | $1.543 \pm 0.043$ | $\mathbf{1 . 4 9 2} \pm \mathbf{0 . 0 5 8}$ | $1.511 \pm 0.029$ |
| mediamill | $(30993,120)$ | - | 1.732 | $1.604 \pm 0.027$ | $1.579 \pm 0.018$ | $\mathbf{1 . 5 5 9} \pm \mathbf{0 . 0 2 2}$ |
| musk | $(7074,168)$ | - | 1.178 | $1.169 \pm 0.006$ | $\mathbf{1 . 1 6 8} \pm \mathbf{0 . 0 0 5}$ | $\mathbf{1 . 1 6 8} \pm \mathbf{0 . 0 0 5}$ |
| CT-slices | $(53500,386)$ | - | 1.242 | $1.240 \pm 0.003$ | $1.235 \pm 0.002$ | $\mathbf{1 . 2 3 4} \pm \mathbf{0 . 0 0 2}$ |
| ISOLET | $(7797,617)$ | - | 1.192 | $1.192 \pm 0.002$ | $\mathbf{1 . 1 8 9} \pm \mathbf{0 . 0 0 1}$ | $\mathbf{1 . 1 8 9} \pm \mathbf{0 . 0 0 0}$ |
| mnist | $(10000,780)$ | - | 1.352 | $1.332 \pm 0.005$ | $1.326 \pm 0.003$ | $\mathbf{1 . 3 2 5} \pm \mathbf{0 . 0 0 3}$ |
| SVHN | $(73257,3072)$ | - | 1.446 | $1.420 \pm 0.005$ | $1.402 \pm 0.009$ | $\mathbf{1 . 3 9 8} \pm \mathbf{0 . 0 0 5}$ |
| ORL | (400, 10304) | - | 1.280 | $1.270 \pm 0.007$ | $1.259 \pm 0.005$ | $\mathbf{1 . 2 5 2} \pm \mathbf{0 . 0 0 4}$ |
| POSS: Count of direct win |  |  |  |  |  |  |
| Average rank |  |  |  |  |  |  |
| $\mathbf{9 . 5}$ | - | $\mathbf{0}$ | $\mathbf{0}$ |  |  |  |
|  |  |  |  |  |  |  |



Figure 1: Error ratio (the smaller, the better) vs. running time of POSS, $\mathrm{PORSS}_{o}$ and $\mathrm{PORSS}_{u}$ on unsupervised feature selection.
are not compared, because POSS has been shown to be better (Qian, Yu, and Zhou 2015; Feng, Qian, and Tang 2019).

As suggested in (Qian, Yu, and Zhou 2015), the number $T$ of iterations of POSS is set to $2 e k^{2} n$. Note that POSS in Algorithm 1 requires one objective evaluation for the newly generated solution $x^{\prime}$ in each iteration, whereas PORSS in Algorithm 2 needs to evaluate two new solutions $\boldsymbol{x}^{\prime \prime}, \boldsymbol{y}^{\prime \prime}$. For the fairness of comparison, the number $T$ of iterations of PORSS is set to $e k^{2} n$; thus, the same number of objective evaluations is used. The budget $k$ is set to 8 . As POSS and PORSS are randomized algorithms, we repeat the running for ten times independently and report the average $f$ values.

Unsupervised Feature Selection. To evaluate a submatrix $\mathbf{S}$, we measure the ratio of its reconstruction error in Definition 3 w.r.t. the smallest rank- $k$ approximation error by SVD:

$$
\text { error ratio }=\left\|\mathbf{A}-\mathbf{S S}^{+} \mathbf{A}\right\|_{F}^{2} /\left\|\mathbf{A}-\mathbf{A}_{k}\right\|_{F}^{2}
$$

where $\mathbf{A}_{k}$ denotes the best rank- $k$ approximation to $\mathbf{A}$ via SVD. The error ratio is larger than 1 , and the smaller the better. The results are summarized in Table 1. Note that the
standard deviation of error ratio is 0 sometimes (e.g., for $\operatorname{PORSS}_{o}$ on the phishing data set), which is because the same good solution is found in ten runs. We can see that the best performance on each data set is always achieved by PORSS $_{o}$ or PORSS $_{u}$. By the sign-test (Demšar 2006) with confidence level 0.05 , POSS is significantly better than the greedy algorithm, consistent with the previous results (Feng, Qian, and Tang 2019), and significantly worse than PORSS ${ }_{o}$ and PORSS ${ }_{u}$, showing the usefulness of recombination. The rank of each algorithm on each data set is also computed as in (Demšar 2006), and averaged in the last row of Table 1.

Sparse Regression. We use $R_{z, S}^{2}$ in Definition 4 to measure the goodness of a subset $S$ of variables. The larger it is, the better. We can see from Table 2 that the algorithms have the similar performance rank as in unsupervised feature selection, i.e., PORSS $_{o}$ and $\operatorname{PORSS}_{u}$ are significantly better than POSS, and the greedy algorithm performs the worst.

To have a more clear comparison, we select the greedy algorithm for the baseline, and plot the curve of error ratio or $R^{2}$ over the running time for POSS, $\mathrm{PORSS}_{o}$ and $\operatorname{PORSS}_{u}$, as shown in Figures 1 and 2. Note that the running time is

Table 2: Sparse regression: the $R^{2}$ value (the larger, the better) of the compared algorithms on ten data sets for $k=8$. The mean $\pm$ std. is reported for randomized algorithms. In each data set, the largest values are bolded. The count of direct win denotes the number of data sets on which POSS has a larger $R^{2}$ value than the corresponding algorithm ( 1 tie is counted as 0.5 win), where significant cells by the sign-test with confidence level 0.05 are bolded.

| Data set | (\#inst, \#feat) | OPT | Greedy | POSS | PORSS $_{o}$ | PORSS $_{u}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| svmguide3 | $(1243,22)$ | 0.221 | 0.214 | $0.220 \pm 0.001$ | $0.220 \pm 0.001$ | $\mathbf{0 . 2 2 1} \pm 0.001$ |
| triazines | $(186,60)$ | 0.328 | 0.316 | $0.327 \pm 0.000$ | $\mathbf{0 . 3 2 8} \pm \mathbf{0 . 0 0 0}$ | $\mathbf{0 . 3 2 8} \pm \mathbf{0 . 0 0 0}$ |
| clean1 | $(476,166)$ | - | 0.371 | $0.386 \pm 0.004$ | $0.387 \pm 0.006$ | $\mathbf{0 . 3 9 3} \pm \mathbf{0 . 0 0 5}$ |
| usps | $(7291,256)$ | - | 0.562 | $0.570 \pm 0.003$ | $\mathbf{0 . 5 7 2} \pm 0.003$ | $\mathbf{0 . 5 7 2} \pm \mathbf{0 . 0 0 3}$ |
| scene | $(1211,294)$ | - | 0.254 | $0.268 \pm 0.003$ | $0.272 \pm 0.002$ | $0.271 \pm 0.002$ |
| protein | $(17766,356)$ | - | 0.132 | $0.132 \pm 0.000$ | $\mathbf{0 . 1 3 3} \pm \mathbf{0 . 0 0 0}$ | $\mathbf{0 . 1 3 3} \pm \mathbf{0 . 0 0 0}$ |
| colon-cancer | $(62,2000)$ | - | 0.890 | $0.906 \pm 0.011$ | $0.909 \pm 0.018$ | $\mathbf{0 . 9 1 1} \pm \mathbf{0 . 0 1 4}$ |
| cifar10 | $(50000,3072)$ | - | 0.069 | $0.070 \pm 0.001$ | $0.070 \pm 0.001$ | $\mathbf{0 . 0 7 1} \pm \mathbf{0 . 0 0 1}$ |
| leukemia | $(72,7129)$ | - | 0.947 | $0.966 \pm 0.009$ | $0.968 \pm 0.006$ | $0.969 \pm 0.007$ |
| smallNORB | (24300, 18432) | - | 0.461 | $0.535 \pm 0.007$ | $0.547 \pm 0.003$ | $\mathbf{0 . 5 5 0} \pm \mathbf{0 . 0 0 2}$ |
| POSS: Count of direct win |  |  | 9.5 | - | 1 | 0 |
| Average rank |  |  | 3.95 | 2.95 | 1.85 | 1.25 |



Figure 2: $R^{2}$ (the larger, the better) vs. running time of POSS, $\mathrm{PORSS}_{o}$ and $\mathrm{PORSS}_{u}$ on sparse regression.
considered in the number of objective function evaluations, and one unit on the $x$-axis corresponds to $k n$ evaluations, the running time of the greedy algorithm. It can be clearly observed that the curves of $\mathrm{PORSS}_{o}$ and $\mathrm{PORSS}_{u}$ are almost always below (above) that of POSS in Figure 1 (Figure 2), implying that $\mathrm{PORSS}_{o}$ and $\mathrm{PORSS}_{u}$ consistently outperform POSS during the running process. It is known that the greedy algorithm is an efficient fixed time algorithm, while PORSS is an anytime algorithm that can use more time to find better solutions. In fact, we find that PORSS can even be both better and faster, e.g., in Figures 1(h-j) and 2(j).

Note that the improvement of PORSS over POSS is small in several cases, which may be because POSS has performed very well. We compute the optimal solution by exhaustive enumeration, denoted as OPT. Due to the computation time limit, OPT is calculated only for the two smallest data sets in both applications. It can be seen from the second and third rows of Tables 1 and 2 that POSS achieves the nearly optimal solution, which also implies that PORSS can bring im-


Figure 3: Comparison for budget $k \in\{6,7,8,9,10\}$.
provement even when POSS has been nearly optimal.
Finally, we examine the influence of budget $k$ in Figure 3. The results for $k \in\{6,7,8,9,10\}$ on the data set CT-slices for unsupervised feature selection and scene for sparse regression show that PORSS always performs the best.

## Conclusion

This paper proposes the PORSS algorithm for subset selection, based on Pareto optimization with recombination. The superiority of PORSS over state-of-the-art algorithms, i.e., POSS and the greedy algorithm, is shown by theoretical analysis, as well as empirical study on the applications of unsupervised feature selection and sparse regression. Theoretical analysis also provides insight on the effect of recombination, which may be helpful for designing improved EAs.

## References

Bäck, T. 1996. Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms. Oxford University Press.
Bhaskara, A.; Rostamizadeh, A.; Altschuler, J.; Zadimoghaddam, M.; Fu, T.; and Mirrokni, V. 2016. Greedy column subset selection: New bounds and distributed algorithms. In Proceedings of the 33 rd International Conference on Machine Learning (ICML'16), 25392548.

Dang, D.-C.; Friedrich, T.; Kötzing, T.; Krejca, M.; Lehre, P. K.; Oliveto, P. S.; Sudholt, D.; and Sutton, A. 2018. Escaping local optima using crossover with emergent diversity. IEEE Transactions on Evolutionary Computation 22(3):484-497.
Das, A., and Kempe, D. 2011. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In Proceedings of the 28th International Conference on Machine Learning (ICML'11), 1057-1064.
Deb, K.; Pratap, A.; Agarwal, S.; and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation 6(2):182-197.
Demšar, J. 2006. Statistical comparisons of classifiers over multiple data sets. Journal of Machine Learning Research 7:1-30.
Doerr, B.; Johannsen, D.; Kötzing, T.; Neumann, F.; and Theile, M. 2013. More effective crossover operators for the all-pairs shortest path problem. Theoretical Computer Science 471:12-26.
Farahat, A. K.; Ghodsi, A.; and Kamel, M. S. 2011. An efficient greedy method for unsupervised feature selection. In Proceedings of the 11th IEEE International Conference on Data Mining (ICDM'11), 161-170.
Feige, U. 1998. A threshold of $\ln n$ for approximating set cover. Journal of the ACM 45(4):634-652.
Feng, C.; Qian, C.; and Tang, K. 2019. Unsupervised feature selection by Pareto optimization. In Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI'19), 3534-3541.
Friedrich, T., and Neumann, F. 2015. Maximizing submodular functions under matroid constraints by evolutionary algorithms. Evolutionary Computation 23(4):543-558.
Giel, O. 2003. Expected runtimes of a simple multi-objective evolutionary algorithm. In Proceedings of the 2003 IEEE Congress on Evolutionary Computation (CEC'03), 1918-1925.
Harshaw, C.; Feldman, M.; Ward, J.; and Karbasi, A. 2019. Submodular maximization beyond non-negativity: Guarantees, fast algorithms, and applications. In Proceedings of the 36th International Conference on Machine Learning (ICML'19), 2634-2643.
Johnson, R. A., and Wichern, D. W. 2007. Applied Multivariate Statistical Analysis. Pearson, 6th edition.
Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'03), 137-146.

Krause, A.; Singh, A.; and Guestrin, C. 2008. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. Journal of Machine Learning Research 9:235284.

Laumanns, M.; Thiele, L.; and Zitzler, E. 2004. Running time analysis of multiobjective evolutionary algorithms on pseudoBoolean functions. IEEE Transactions on Evolutionary Computation 8(2):170-182.

Lin, H., and Bilmes, J. 2011. A class of submodular functions for document summarization. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL'11), 510-520.
Miller, A. 2002. Subset Selection in Regression. Chapman and Hall/CRC, 2nd edition.
Nemhauser, G. L., and Wolsey, L. A. 1978. Best algorithms for approximating the maximum of a submodular set function. Mathematics of Operations Research 3(3):177-188.
Nemhauser, G. L.; Wolsey, L. A.; and Fisher, M. L. 1978. An analysis of approximations for maximizing submodular set functions I. Mathematical Programming 14(1):265-294.

Neumann, F., and Theile, M. 2010. How crossover speeds up evolutionary algorithms for the multi-criteria all-pairs-shortest-path problem. In Proceedings of the 11th International Conference on Parallel Problem Solving from Nature (PPSN'10), 667-676.
Oliveto, P. S., and Witt, C. 2014. On the runtime analysis of the simple genetic algorithm. Theoretical Computer Science 545:2-19.
Ordozgoiti, B.; Canaval, S.; and Mozo, A. 2018. Iterative column subset selection. Knowledge and Information Systems 54(1):6594.

Qian, C.; Shi, J.-C.; Yu, Y.; Tang, K.; and Zhou, Z.-H. 2016. Parallel Pareto optimization for subset selection. In Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI'16), 1939-1945.
Qian, C.; Shi, J.-C.; Yu, Y.; Tang, K.; and Zhou, Z.-H. 2017. Subset selection under noise. In Advances in Neural Information Processing Systems 30 (NIPS'17), 3562-3572.
Qian, C.; Li, G.; Feng, C.; and Tang, K. 2018. Distributed Pareto optimization for subset selection. In Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18), 1492-1498.
Qian, C.; Yu, Y.; and Zhou, Z.-H. 2013. An analysis on recombination in multi-objective evolutionary optimization. Artificial Intelligence 204:99-119.
Qian, C.; Yu, Y.; and Zhou, Z.-H. 2015. Subset selection by Pareto optimization. In Advances in Neural Information Processing Systems 28 (NIPS'15), 1765-1773.
Qian, C. 2019. Distributed Pareto optimization for large-scale noisy subset selection. IEEE Transactions on Evolutionary Computation in press.
Roostapour, V.; Neumann, A.; Neumann, F.; and Friedrich, T. 2019. Pareto optimization for subset selection with dynamic cost constraints. In Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI'19), 2354-2361.
Sudholt, D. 2017. How crossover speeds up building block assembly in genetic algorithms. Evolutionary Computation 25(2):237274.

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Methodological) 58(1):267-288.


[^0]:    *This work was supported by the NSFC (61603367).
    Copyright (C) 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ https://archive.ics.uci.edu/ml/datasets.html, https://www.csie. ntu.edu.tw/~cjlin/libsvmtools/datasets/ and http://www.cl.cam.ac. $\mathrm{uk} /$ research/dtg/attarchive/facedatabase.html.

