


Best of Both Worlds: Practical and Theoretically Optimal Submodular Maximization in Parallel

Yixin Chen, Tonmoy Dey, Alan Kuhnle
Department of Computer Science
Florida State University

Submodularity

Can be defined as a property of a function where:

- Given an objective function $f(S) = \left| \bigcup_{s \in S} \text{area}(s) \right|$




$$A = [S1, S2]$$

$$f(A \cup S^*) - f(A)$$


Submodularity

Can be defined as a property of a function where:

- Given an objective function $f(S) = \left| \bigcup_{s \in S} \text{area}(s) \right|$
- The marginal gain of adding an element to a set **diminishes with increase in size of the set**



$A = [S1, S2]$



$B = [S1, S2, S3]$

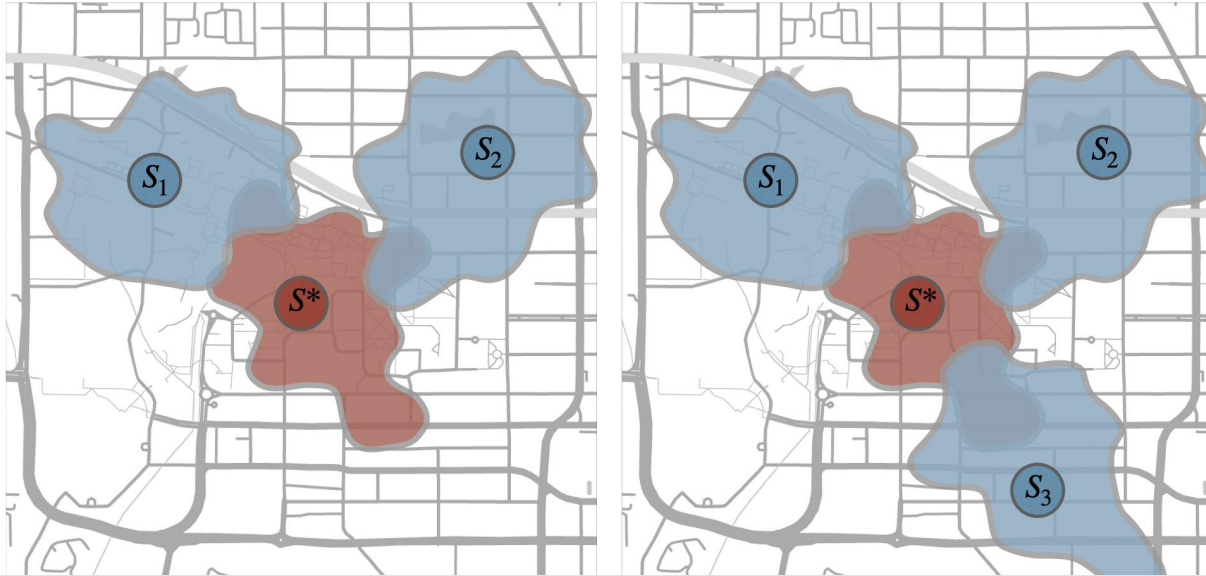
$$f(A \cup S^*) - f(A) \geq f(B \cup S^*) - f(B)$$

Monotonicity

A function $f: 2^V \rightarrow \mathbb{R}$ is monotone if:

- for every $A \subseteq B \subseteq V, f(A) \leq f(B)$
- Or, for every $A \subseteq V$ and $e \in V$
- it holds that $\Delta(e | A) \geq 0$

Submodular Maximization - Cardinality Constraint



Given a cardinality constraint k

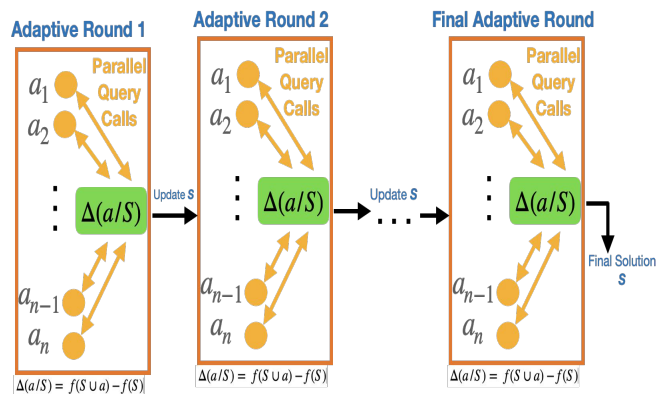
Objective: Maximize the coverage function with no more than k elements:

i.e. Maximize $f(S)$, subject to $|S| \leq k$

Submodular Maximization - Performance Metrics

- **Metrics**

- Approximation Ratio:
 - the minimal ratio of the solution to the optimal result
- Query Complexity:
 - total number of query calls
- Adaptivity:
 - Introduced by Balkanski and Singer¹ for submodular optimization
 - Defined as the minimal number of sequential rounds required to achieve a constant factor approximation when polynomially-many queries can be executed in parallel at each round.
 - It is the metric used to define how efficiently the algorithm can parallelize each iteration



Related Work

- Optimal ratio¹: $1 - 1/e$
- Lower bound of query complexity²: $\Omega(n)$
- Lower bound of adaptivity³: $\Omega(\log(n)/\log \log(n))$

[1] G L Nemhauser and L A Wolsey. **Best Algorithms for Approximating the Maximum of a Submodular Set Function**. Mathematics of Operations Research, 3(3):177–188, 1978.

[2] Alan Kuhnle. **Quick Streaming Algorithms for Maximization of Monotone Submodular Functions in Linear Time**. In Artificial Intelligence and Statistics (AISTATS), 2021.

[3] Eric Balkanski and Yaron Singer. **The adaptive complexity of maximizing a submodular function**. In ACM SIGACT Symposium on Theory of Computing (STOC), 2018.

Related Work

- Several previous works get nearly theoretically optimal result
- Impractical with large constant factors

Reference	Ratio	Adaptivity	Queries
Ene and Nguyen [14]	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O(n \text{poly}(\log n, 1/\varepsilon))$
Chekuri and Quanrud [11] (RPG)	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right) \dagger$	$O\left(\frac{n}{\varepsilon^4} \log(n)\right) \dagger$
Fahrbach et al. [17] (BSM)	$1 - 1/e - \varepsilon \dagger$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O\left(\frac{n}{\varepsilon^3} \log \log k\right) \dagger$
Fahrbach et al. [17] (SM)	$1 - 1/e - \varepsilon \dagger$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O\left(\frac{n}{\varepsilon^3} \log(1/\varepsilon)\right) \dagger$
Breuer et al. [9] (FAST)	$1 - 1/e - \varepsilon \dagger \ddagger$	$O\left(\frac{1}{\varepsilon^2} \log(n) \log^2\left(\frac{\log(k)}{\varepsilon}\right)\right)$	$O\left(\frac{n}{\varepsilon^2} \log\left(\frac{\log(k)}{\varepsilon}\right)\right)$
LS+PGB [Theorem 3]	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n/\varepsilon)\right)$	$O\left(\frac{n}{\varepsilon^2}\right) \dagger$

\dagger indicates the result holds with constant probability or in expectation;

\ddagger indicates the result does not hold on all instances of SM;

while no symbol indicates the result holds with probability greater than $1-O(1/n)$

Related Work - FAST¹

- speed up the algorithms using the adaptive sequencing technique
- Sacrifice the theoretical guarantee
- Significantly, no ratio for $k < 850$

Reference	Ratio	Adaptivity	Queries
Ene and Nguyen [14]	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O(n \text{poly}(\log n, 1/\varepsilon))$
Chekuri and Quanrud [11] (RPG)	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right) \dagger$	$O\left(\frac{n}{\varepsilon^4} \log(n)\right) \dagger$
Fahrbach et al. [17] (BSM)	$1 - 1/e - \varepsilon \dagger$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O\left(\frac{n}{\varepsilon^3} \log \log k\right) \dagger$
Fahrbach et al. [17] (SM)	$1 - 1/e - \varepsilon \dagger$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O\left(\frac{n}{\varepsilon^3} \log(1/\varepsilon)\right) \dagger$
Breuer et al. [9] (FAST)	$1 - 1/e - \varepsilon \dagger \ddagger$	$O\left(\frac{1}{\varepsilon^2} \log(n) \log^2\left(\frac{\log(k)}{\varepsilon}\right)\right)$	$O\left(\frac{n}{\varepsilon^2} \log\left(\frac{\log(k)}{\varepsilon}\right)\right)$
LS+PGB [Theorem 3]	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n/\varepsilon)\right)$	$O\left(\frac{n}{\varepsilon^2}\right) \dagger$

\dagger indicates the result holds with constant probability or in expectation;

\ddagger indicates the result does not hold on all instances of SM;

while no symbol indicates the result holds with probability greater than $1 - O(1/n)$

[1] Adam Breuer, Eric Balkanski, and Yaron Singer. The FAST Algorithm for Submodular Maximization. In International Conference on Machine Learning (ICML), 2019.

Our Main Algorithm - LS+PGB

- Provides theoretical guarantee for all k values
- Empirically outperforms, in terms of runtime, adaptivity, total queries, and objective values, the previous state-of-the-art algorithm FAST (Breuer et al.)

Reference	Ratio	Adaptivity	Queries
Ene and Nguyen [14]	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O(n \text{poly}(\log n, 1/\varepsilon))$
Chekuri and Quanrud [11] (RPG)	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right) \dagger$	$O\left(\frac{n}{\varepsilon^4} \log(n)\right) \dagger$
Fahrbach et al. [17] (BSM)	$1 - 1/e - \varepsilon \dagger$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O\left(\frac{n}{\varepsilon^3} \log \log k\right) \dagger$
Fahrbach et al. [17] (SM)	$1 - 1/e - \varepsilon \dagger$	$O\left(\frac{1}{\varepsilon^2} \log(n)\right)$	$O\left(\frac{n}{\varepsilon^3} \log(1/\varepsilon)\right) \dagger$
Breuer et al. [9] (FAST)	$1 - 1/e - \varepsilon \dagger \ddagger$	$O\left(\frac{1}{\varepsilon^2} \log(n) \log^2\left(\frac{\log(k)}{\varepsilon}\right)\right)$	$O\left(\frac{n}{\varepsilon^2} \log\left(\frac{\log(k)}{\varepsilon}\right)\right)$
LS+PGB [Theorem 3]	$1 - 1/e - \varepsilon$	$O\left(\frac{1}{\varepsilon^2} \log(n/\varepsilon)\right)$	$O\left(\frac{n}{\varepsilon^2}\right) \dagger$

\dagger indicates the result holds with constant probability or in expectation;

\ddagger indicates the result does not hold on all instances of SM;

while no symbol indicates the result holds with probability greater than $1-O(1/n)$

Contributions

- **LinearSeq**: obtains a constant ratio $(4 + O(\varepsilon))^{-1}$ in expected $O(n)$ query complexity and $O(\log(n))$ adaptivity with probability $1 - 1/n$
 - Modified: achieves $O(\log(n/k))$ adaptivity with sacrificing the ratio to be $(5 + O(\varepsilon))^{-1}$

Contributions

- **LinearSeq**: obtains a constant factor $(4 + O(\varepsilon))^{-1}$ in expected $O(n)$ query complexity and $O(\log(n))$ adaptivity with probability $1 - 1/n$
 - Modified: achieves $O(\log(n/k))$ adaptivity with sacrificing the ratio to be $(5 + O(\varepsilon))^{-1}$
- **ThresholdSeq**: the average marginal gain is larger than a specified threshold with probability $1 - 1/n$ in expected $O(n)$ query complexity and adaptive rounds $O(\log(n))$

Contributions

- **LinearSeq**: obtains a constant factor $(4 + O(\varepsilon))^{-1}$ in expected $O(n)$ query complexity and $O(\log(n))$ adaptivity with probability $1 - 1/n$
 - Modified: achieves $O(\log(n/k))$ adaptivity with sacrificing the ratio to be $(5 + O(\varepsilon))^{-1}$
- **ThresholdSeq**: the average marginal gain is larger than a specified threshold with probability $1 - 1/n$ in expected $O(n)$ query complexity and adaptive rounds $O(\log(n))$
- **LS+PGB**: uses LinearSeq as preprocessing algorithm and combines ThresholdSeq with boost mechanism; obtains nearly the optimal result

Highly adaptive linear-time algorithm - $\frac{1}{4}$ ratio

Algorithm 2 Highly Adaptive Linear-Time Algorithm

```
1: Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ 
2: Initialize  $A \leftarrow \emptyset$ 
3: for  $u \in \mathcal{N}$  do
4:   if  $\Delta(u|A) \geq f(A)/k$  then
5:      $A \leftarrow A \cup \{u\}$ 
6: return  $A' \leftarrow \{\text{last } k \text{ elements added to } A\}$ 
```

- Use $f(A)/k$ as threshold to ensure that:
 - The last k elements in A contain a constant fraction of the value $f(A)$
 - $f(A)$ is within a constant fraction of OPT
- Totally $2n$ query calls and n adaptive rounds

Highly adaptive linear-time algorithm - $\frac{1}{4}$ ratio

Algorithm 2 Highly Adaptive Linear-Time Algorithm

```
1: Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ 
2: Initialize  $A \leftarrow \emptyset$ 
3: for  $u \in \mathcal{N}$  do
4:   if  $\Delta(u|A) \geq f(A)/k$  then
5:      $A \leftarrow A \cup \{u\}$ 
6: return  $A' \leftarrow \{\text{last } k \text{ elements added to } A\}$ 
```

Problem: How to parallelize this algorithm to a lowly adaptive version without loss much of approximation ratio and query complexity

- Use $f(A)/k$ as threshold to ensure that:
 - The last k elements in A contain a constant fraction of the value $f(A)$
 - $f(A)$ is within a constant fraction of OPT
- Totally $2n$ query calls and n adaptive rounds

LinearSeq

- $1/4$ ratio, $O(n)$ query and $O(\log n)$ adaptive complexity

Algorithm 1 The algorithm that obtains ratio $(4 + O(\varepsilon))^{-1}$ in $O(\log(n)/\varepsilon^3)$ adaptive rounds and expected $O(n/\varepsilon^3)$ queries.

```
1: procedure LINEARSEQ( $f, \mathcal{N}, k, \varepsilon$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ , error  $\varepsilon$ 
3:    $a = \arg \max_{u \in \mathcal{N}} f(\{u\})$ 
4:   Initialize  $A \leftarrow \{a\}$ ,  $V \leftarrow \mathcal{N}$ ,  $\ell = \lceil 4(1 + 1/(\beta\varepsilon)) \log(n) \rceil$ ,  $\beta = \varepsilon/(16 \log(8/(1 - e^{-\varepsilon/2})))$ 
5:   for  $j \leftarrow 1$  to  $\ell$  do
6:     Update  $V \leftarrow \{x \in V : \Delta(x | A) \geq f(A)/k\}$  and filter out the rest
7:     if  $|V| = 0$  then break
8:      $V = \{v_1, v_2, \dots, v_{|V|}\} \leftarrow \text{random-permutation}(V)$ 
9:      $\Lambda \leftarrow \{ \lfloor (1 + \varepsilon)^u \rfloor : 1 \leq \lfloor (1 + \varepsilon)^u \rfloor \leq k, u \in \mathbb{N} \}$ 
            $\cup \{ \lfloor k + u\varepsilon k \rfloor : \lfloor k + u\varepsilon k \rfloor \leq |V|, u \in \mathbb{N} \} \cup \{|V|\}$ 
10:     $B[\lambda_i] = \text{false}$ , for  $\lambda_i \in \Lambda$ 
11:    for  $\lambda_i \in \Lambda$  in parallel do
12:       $T_{\lambda_{i-1}} \leftarrow \{v_1, v_2, \dots, v_{\lambda_{i-1}}\}$ ;  $T_{\lambda_i} \leftarrow \{v_1, v_2, \dots, v_{\lambda_i}\}$ ;  $T'_{\lambda_i} \leftarrow T_{\lambda_i} \setminus T_{\lambda_{i-1}}$ 
13:      if  $\Delta(T'_{\lambda_i} | A \cup T_{\lambda_{i-1}}) / |T'_{\lambda_i}| \geq (1 - \varepsilon)f(A \cup T_{\lambda_{i-1}})/k$  then  $B[\lambda_i] \leftarrow \text{true}$ 
14:       $\lambda^* \leftarrow \max\{\lambda_i \in \Lambda : B[\lambda_i] = \text{false} \text{ and } ((\lambda_i \leq k \text{ and } B[1] \text{ to } B[\lambda_{i-1}] \text{ are all true}) \text{ or } (\lambda_i > k \text{ and } \exists m \geq 1 \text{ s.t. } |\bigcup_{u=m}^{i-1} T'_{\lambda_u}| \geq k \text{ and } B[\lambda_m] \text{ to } B[\lambda_{i-1}] \text{ are all true}))\}$ 
15:       $A \leftarrow A \cup T_{\lambda^*}$ 
16:    if  $|V| > 0$  then return failure
17:    return  $A' \leftarrow$  last  $k$  elements added to  $A$ 
```

- filter out the elements with small marginal gains

LinearSeq

- $1/4$ ratio, $O(n)$ query and $O(\log n)$ adaptive complexity

Algorithm 1 The algorithm that obtains ratio $(4 + O(\varepsilon))^{-1}$ in $O(\log(n)/\varepsilon^3)$ adaptive rounds and expected $O(n/\varepsilon^3)$ queries.

```
1: procedure LINEARSEQ( $f, \mathcal{N}, k, \varepsilon$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ , error  $\varepsilon$ 
3:    $a = \arg \max_{u \in \mathcal{N}} f(\{u\})$ 
4:   Initialize  $A \leftarrow \{a\}$ ,  $V \leftarrow \mathcal{N}$ ,  $\ell = \lceil 4(1 + 1/(\beta\varepsilon)) \log(n) \rceil$ ,  $\beta = \varepsilon/(16 \log(8/(1 - e^{-\varepsilon/2})))$ 
5:   for  $j \leftarrow 1$  to  $\ell$  do
6:     Update  $V \leftarrow \{x \in V : \Delta(x | A) \geq f(A)/k\}$  and filter out the rest
7:     if  $|V| = 0$  then break
8:      $V = \{v_1, v_2, \dots, v_{|V|}\} \leftarrow \text{random-permutation}(V)$ 
9:      $\Lambda \leftarrow \{ \lfloor (1 + \varepsilon)^u \rfloor : 1 \leq \lfloor (1 + \varepsilon)^u \rfloor \leq k, u \in \mathbb{N} \}$ 
            $\cup \{ \lfloor k + u\varepsilon k \rfloor : \lfloor k + u\varepsilon k \rfloor \leq |V|, u \in \mathbb{N} \} \cup \{|V|\}$ 
10:     $B[\lambda_i] = \text{false}$ , for  $\lambda_i \in \Lambda$ 
11:    for  $\lambda_i \in \Lambda$  in parallel do
12:       $T_{\lambda_{i-1}} \leftarrow \{v_1, v_2, \dots, v_{\lambda_{i-1}}\}$ ;  $T_{\lambda_i} \leftarrow \{v_1, v_2, \dots, v_{\lambda_i}\}$ ;  $T'_{\lambda_i} \leftarrow T_{\lambda_i} \setminus T_{\lambda_{i-1}}$ 
13:      if  $\Delta(T'_{\lambda_i} | A \cup T_{\lambda_{i-1}}) / |T'_{\lambda_i}| \geq (1 - \varepsilon)f(A \cup T_{\lambda_{i-1}})/k$  then  $B[\lambda_i] \leftarrow \text{true}$ 
14:       $\lambda^* \leftarrow \max\{\lambda_i \in \Lambda : B[\lambda_i] = \text{false} \text{ and } ((\lambda_i \leq k \text{ and } B[1] \text{ to } B[\lambda_{i-1}] \text{ are all true}) \text{ or } (\lambda_i > k \text{ and } \exists m \geq 1 \text{ s.t. } |\bigcup_{u=m}^{i-1} T'_{\lambda_u}| \geq k \text{ and } B[\lambda_m] \text{ to } B[\lambda_{i-1}] \text{ are all true}))\}$ 
15:       $A \leftarrow A \cup T_{\lambda^*}$ 
16:    if  $|V| > 0$  then return failure
17:    return  $A' \leftarrow$  last  $k$  elements added to  $A$ 
```

- get a randomly selected sequence

LinearSeq

- $1/4$ ratio, $O(n)$ query and $O(\log n)$ adaptive complexity

Algorithm 1 The algorithm that obtains ratio $(4 + O(\varepsilon))^{-1}$ in $O(\log(n)/\varepsilon^3)$ adaptive rounds and expected $O(n/\varepsilon^3)$ queries.

```
1: procedure LINEARSEQ( $f, \mathcal{N}, k, \varepsilon$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ , error  $\varepsilon$ 
3:    $a = \arg \max_{u \in \mathcal{N}} f(\{u\})$ 
4:   Initialize  $A \leftarrow \{a\}$ ,  $V \leftarrow \mathcal{N}$ ,  $\ell = \lceil 4(1 + 1/(\beta\varepsilon)) \log(n) \rceil$ ,  $\beta = \varepsilon/(16 \log(8/(1 - e^{-\varepsilon/2})))$ 
5:   for  $j \leftarrow 1$  to  $\ell$  do
6:     Update  $V \leftarrow \{x \in V : \Delta(x | A) \geq f(A)/k\}$  and filter out the rest
7:     if  $|V| = 0$  then break
8:      $V = \{v_1, v_2, \dots, v_{|V|}\} \leftarrow \mathbf{random-permutation}(V)$ 
9:      $\Lambda \leftarrow \{ \lfloor (1 + \varepsilon)^u \rfloor : 1 \leq \lfloor (1 + \varepsilon)^u \rfloor \leq k, u \in \mathbb{N} \}$   

        $\cup \{ \lfloor k + u\varepsilon k \rfloor : \lfloor k + u\varepsilon k \rfloor \leq |V|, u \in \mathbb{N} \} \cup \{|V|\}$ 
10:     $B[\lambda_i] = \mathbf{false}$ , for  $\lambda_i \in \Lambda$ 
11:    for  $\lambda_i \in \Lambda$  in parallel do
12:       $T_{\lambda_{i-1}} \leftarrow \{v_1, v_2, \dots, v_{\lambda_{i-1}}\}$ ;  $T_{\lambda_i} \leftarrow \{v_1, v_2, \dots, v_{\lambda_i}\}$ ;  $T'_{\lambda_i} \leftarrow T_{\lambda_i} \setminus T_{\lambda_{i-1}}$ 
13:      if  $\Delta(T'_{\lambda_i} | A \cup T_{\lambda_{i-1}}) / |T'_{\lambda_i}| \geq (1 - \varepsilon)f(A \cup T_{\lambda_{i-1}})/k$  then  $B[\lambda_i] \leftarrow \mathbf{true}$ 
14:       $\lambda^* \leftarrow \max\{\lambda_i \in \Lambda : B[\lambda_i] = \mathbf{false} \text{ and } ((\lambda_i \leq k \text{ and } B[1] \text{ to } B[\lambda_{i-1}] \text{ are all true}) \text{ or } (\lambda_i > k \text{ and } \exists m \geq 1 \text{ s.t. } |\bigcup_{u=m}^{i-1} T'_{\lambda_u}| \geq k \text{ and } B[\lambda_m] \text{ to } B[\lambda_{i-1}] \text{ are all true}))\}$ 
15:       $A \leftarrow A \cup T_{\lambda^*}$ 
16:    if  $|V| > 0$  then return failure
17:    return  $A' \leftarrow$  last  $k$  elements added to  $A$ 
```

- split V into blocks to reduce the query calls

LinearSeq

- $1/4$ ratio, $O(n)$ query and $O(\log n)$ adaptive complexity

Algorithm 1 The algorithm that obtains ratio $(4 + O(\varepsilon))^{-1}$ in $O(\log(n)/\varepsilon^3)$ adaptive rounds and expected $O(n/\varepsilon^3)$ queries.

```
1: procedure LINEARSEQ( $f, \mathcal{N}, k, \varepsilon$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ , error  $\varepsilon$ 
3:    $a = \arg \max_{u \in \mathcal{N}} f(\{u\})$ 
4:   Initialize  $A \leftarrow \{a\}$ ,  $V \leftarrow \mathcal{N}$ ,  $\ell = \lceil 4(1 + 1/(\beta\varepsilon)) \log(n) \rceil$ ,  $\beta = \varepsilon/(16 \log(8/(1 - e^{-\varepsilon/2})))$ 
5:   for  $j \leftarrow 1$  to  $\ell$  do
6:     Update  $V \leftarrow \{x \in V : \Delta(x | A) \geq f(A)/k\}$  and filter out the rest
7:     if  $|V| = 0$  then break
8:      $V = \{v_1, v_2, \dots, v_{|V|}\} \leftarrow \text{random-permutation}(V)$ 
9:      $\Lambda \leftarrow \{ \lfloor (1 + \varepsilon)^u \rfloor : 1 \leq \lfloor (1 + \varepsilon)^u \rfloor \leq k, u \in \mathbb{N} \}$ 
            $\cup \{ \lfloor k + u\varepsilon k \rfloor : \lfloor k + u\varepsilon k \rfloor \leq |V|, u \in \mathbb{N} \} \cup \{|V|\}$ 
10:     $B[\lambda_i] = \text{false}$ , for  $\lambda_i \in \Lambda$ 
11:    for  $\lambda_i \in \Lambda$  in parallel do
12:       $T_{\lambda_{i-1}} \leftarrow \{v_1, v_2, \dots, v_{\lambda_{i-1}}\}$ ;  $T_{\lambda_i} \leftarrow \{v_1, v_2, \dots, v_{\lambda_i}\}$ ;  $T'_{\lambda_i} \leftarrow T_{\lambda_i} \setminus T_{\lambda_{i-1}}$ 
13:      if  $\Delta(T'_{\lambda_i} | A \cup T_{\lambda_{i-1}}) / |T'_{\lambda_i}| \geq (1 - \varepsilon)f(A \cup T_{\lambda_{i-1}})/k$  then  $B[\lambda_i] \leftarrow \text{true}$ 
14:       $\lambda^* \leftarrow \max\{\lambda_i \in \Lambda : B[\lambda_i] = \text{false} \text{ and } ((\lambda_i \leq k \text{ and } B[1] \text{ to } B[\lambda_{i-1}] \text{ are all true}) \text{ or } (\lambda_i > k \text{ and } \exists m \geq 1 \text{ s.t. } |\bigcup_{u=m}^{i-1} T'_{\lambda_u}| \geq k \text{ and } B[\lambda_m] \text{ to } B[\lambda_{i-1}] \text{ are all true}))\}$ 
15:       $A \leftarrow A \cup T_{\lambda^*}$ 
16:    if  $|V| > 0$  then return failure
17:    return  $A' \leftarrow$  last  $k$  elements added to  $A$ 
```

- check the marginal gain of each block in parallel

LinearSeq

- $1/4$ ratio, $O(n)$ query and $O(\log n)$ adaptive complexity

Algorithm 1 The algorithm that obtains ratio $(4 + O(\varepsilon))^{-1}$ in $O(\log(n)/\varepsilon^3)$ adaptive rounds and expected $O(n/\varepsilon^3)$ queries.

```
1: procedure LINEARSEQ( $f, \mathcal{N}, k, \varepsilon$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ , error  $\varepsilon$ 
3:    $a = \arg \max_{u \in \mathcal{N}} f(\{u\})$ 
4:   Initialize  $A \leftarrow \{a\}$ ,  $V \leftarrow \mathcal{N}$ ,  $\ell = \lceil 4(1 + 1/(\beta\varepsilon)) \log(n) \rceil$ ,  $\beta = \varepsilon/(16 \log(8/(1 - e^{-\varepsilon/2})))$ 
5:   for  $j \leftarrow 1$  to  $\ell$  do
6:     Update  $V \leftarrow \{x \in V : \Delta(x | A) \geq f(A)/k\}$  and filter out the rest
7:     if  $|V| = 0$  then break
8:      $V = \{v_1, v_2, \dots, v_{|V|}\} \leftarrow \text{random-permutation}(V)$ 
9:      $\Lambda \leftarrow \{ \lfloor (1 + \varepsilon)^u \rfloor : 1 \leq \lfloor (1 + \varepsilon)^u \rfloor \leq k, u \in \mathbb{N} \}$ 
            $\cup \{ \lfloor k + u\varepsilon k \rfloor : \lfloor k + u\varepsilon k \rfloor \leq |V|, u \in \mathbb{N} \} \cup \{|V|\}$ 
10:     $B[\lambda_i] = \text{false}$ , for  $\lambda_i \in \Lambda$ 
11:    for  $\lambda_i \in \Lambda$  in parallel do
12:       $T_{\lambda_{i-1}} \leftarrow \{v_1, v_2, \dots, v_{\lambda_{i-1}}\}$ ;  $T_{\lambda_i} \leftarrow \{v_1, v_2, \dots, v_{\lambda_i}\}$ ;  $T'_{\lambda_i} \leftarrow T_{\lambda_i} \setminus T_{\lambda_{i-1}}$ 
13:      if  $\Delta(T'_{\lambda_i} | A \cup T_{\lambda_{i-1}}) / |T'_{\lambda_i}| \geq (1 - \varepsilon)f(A \cup T_{\lambda_{i-1}})/k$  then  $B[\lambda_i] \leftarrow \text{true}$ 
14:       $\lambda^* \leftarrow \max\{\lambda_i \in \Lambda : B[\lambda_i] = \text{false} \text{ and } ((\lambda_i \leq k \text{ and } B[1] \text{ to } B[\lambda_{i-1}] \text{ are all true}) \text{ or } (\lambda_i > k \text{ and } \exists m \geq 1 \text{ s.t. } |\bigcup_{u=m}^{i-1} T'_{\lambda_u}| \geq k \text{ and } B[\lambda_m] \text{ to } B[\lambda_{i-1}] \text{ are all true}))\}$ 
15:       $A \leftarrow A \cup T_{\lambda^*}$ 
16:    if  $|V| > 0$  then return failure
17:    return  $A' \leftarrow$  last  $k$  elements added to  $A$ 
```

- prefix selection
 - ensure that the selected subset obtains large marginal gain
 - ensure that an $\varepsilon/2$ -fraction of V can be filtered out at next iteration with high probability

ThresholdSeq

- $O(n)$ query and $O(\log n)$ adaptive complexity

Algorithm 3 A Parallelizable Greedy Algorithm for Fixed Threshold τ

```
1: procedure THRESHOLDSEQ( $f, \mathcal{N}, k, \delta, \varepsilon, \tau$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ , revision  $\delta$ , error  $\varepsilon$ , threshold  $\tau$ 
3:   Initialize  $A \leftarrow \emptyset$ ,  $V \leftarrow \mathcal{N}$ ,  $\ell = \lceil 4(1 + 2/\varepsilon) \log(n/\delta) \rceil$ 
4:   for  $j \leftarrow 1$  to  $\ell$  do
5:     Update  $V \leftarrow \{x \in V : \Delta(x | A) \geq \tau\}$  and filter out the rest
6:     if  $|V| = 0$  then
7:       return  $A$ 
8:      $V \leftarrow \text{random-permutation}(V)$ .
9:      $s \leftarrow \min\{k - |A|, |V|\}$ 
10:     $\Lambda \leftarrow \{\lfloor (1 + \varepsilon)^u \rfloor : 1 \leq \lfloor (1 + \varepsilon)^u \rfloor \leq s, u \in \mathbb{N}\} \cup \{s\}$ 
11:     $B \leftarrow \emptyset$ 
12:    for  $\lambda_i \in \Lambda$  in parallel do
13:       $T_{\lambda_i} \leftarrow \{v_1, v_2, \dots, v_{\lambda_i}\}$ 
14:      if  $\Delta(T_{\lambda_i} | A) / |T_{\lambda_i}| \geq (1 - \varepsilon)\tau$  then
15:         $B \leftarrow B \cup \{\lambda_i\}$ 
16:     $\lambda^* \leftarrow \min\{\lambda_i \in \Lambda : \lambda_i > b, \forall b \in B\}$ 
17:     $A \leftarrow A \cup T_{\lambda^*}$ 
18:    if  $|A| = k$  then
19:      return  $A$ 
20: return failure
```

- analogous to LinearSeq
- $f(A)/|A| \geq (1 - \varepsilon)\tau / (1 + \varepsilon)$
- constant threshold τ
- much simpler while only consider the first $\min\{k - |A|, |V|\}$ elements
- stop when $|A|=k$

LS+PGB

- $1-1/e$ ratio, $O(n)$ query and $O(\log n)$ adaptive complexity

Algorithm 4 A Parallelizable Greedy Algorithm to Boost to the Optimal Ratio.

```
1: procedure PARALLELGREEDYBOOST( $f, \mathcal{N}, k, \alpha, \Gamma, \varepsilon$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , constraint  $k$ , constant  $\alpha$ , value  $\Gamma$  such that  $\Gamma \leq f(O) \leq \Gamma/\alpha$ , error  $\varepsilon$ 
3:   Initialize  $\tau \leftarrow \Gamma/(\alpha k)$ ,  $\delta \leftarrow 1/(\log_{1-\varepsilon}(\alpha/3) + 1)$ ,  $A \leftarrow \emptyset$ 
4:   while  $\tau \geq \Gamma/(3k)$  do
5:      $\tau \leftarrow \tau(1 - \varepsilon)$ 
6:      $S \leftarrow \text{THRESHOLDSEQ}(f_A, \mathcal{N}, k - |A|, \delta, \varepsilon/3, \tau)$ 
7:      $A \leftarrow A \cup S$ 
8:     if  $|A| = k$  then
9:       return  $A$ 
10:  return  $A$ 
```

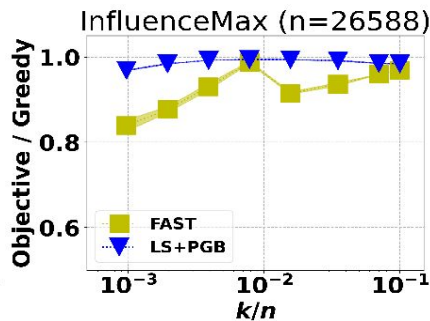

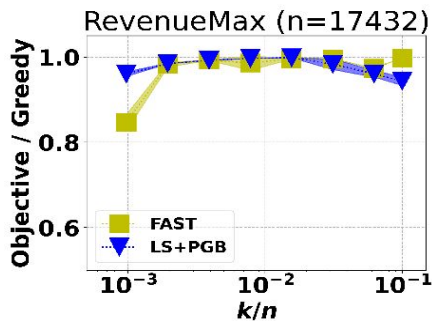
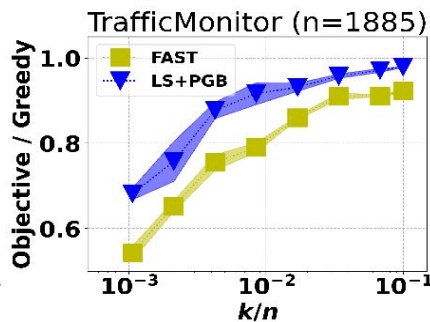


- Use LinearSeq as preprocessing algorithm to get an α -approximation solution I
- Γ and α are used to produce an initial threshold value τ
- The threshold value is iteratively decreased by the factor $(1-\varepsilon)$

Empirical Results - Environment Setup

- The experiments are conducted on a server running **Ubuntu 20.04.2** with kernel **5.8.0**
- The hardware of the system consists of **40 Intel(R) Xeon(R) Gold 5218R CPU @ 2.10GHz** cores with **75 threads** made available to the algorithms for the experiments.
- All algorithms were implemented using **Open-MPI** and **mpi4py** library
- The experiments were performed across **six applications** with **groundset size ranging from 1,885 - 100,000** and **K values ranging from 0.1% - 10% of groundset**.
- For evaluation, the metrics of **total time**, **total queries**, **adaptive rounds** and **objective value** were used for comparison.
- Our algorithm (**LS+PGB**) is compared to the previous state-of-art algorithm **FAST¹**.

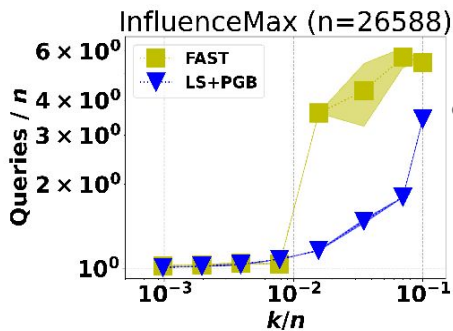
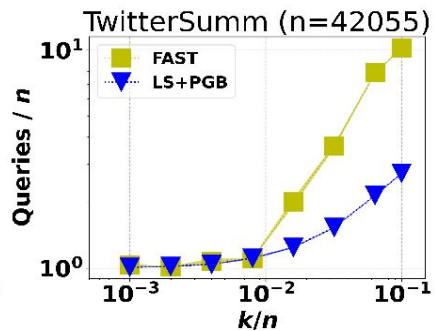
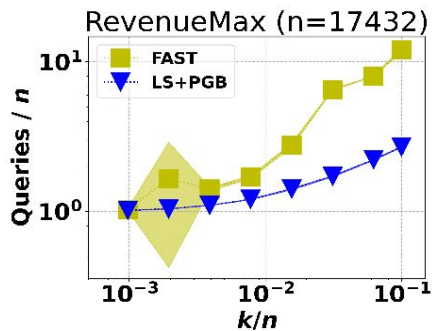



[1] Adam Breuer, Eric Balkanski, and Yaron Singer. The FAST Algorithm for Submodular Maximization. In International Conference on Machine Learning (ICML), 2019.

Empirical Results - Objective Value



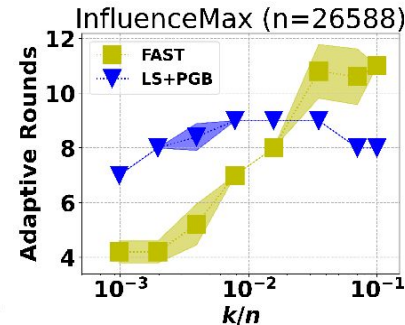
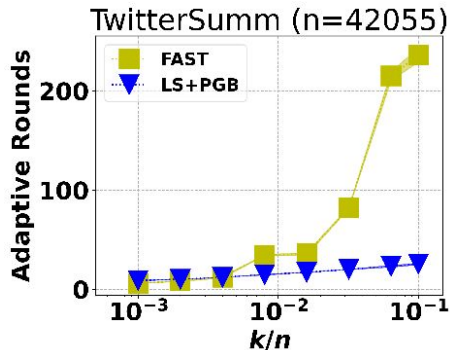
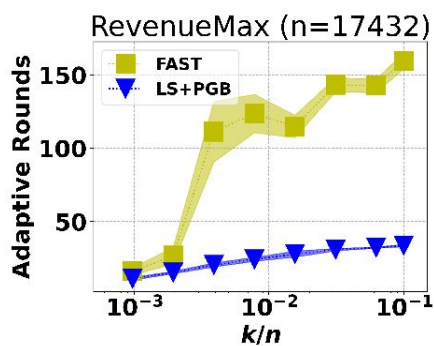
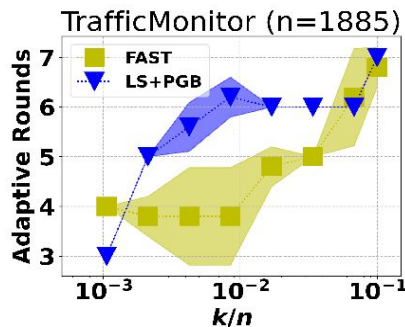
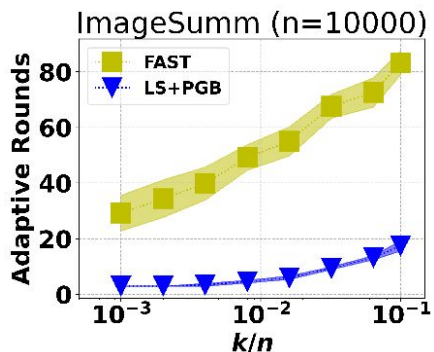
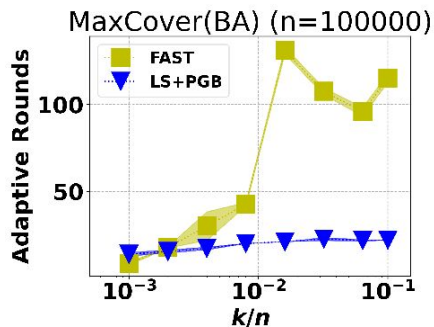
- The objective value is normalized by that of **Greedy**
- Overall **LS+PGB** either maintains or outperforms the objective obtained by **FAST** across all applications
- With the **TrafficMonitor** and **MaxCover (BA)** being the instances where it **exceeds the average objective value of FAST by 6% and 5% respectively.**

Empirical Results - Queries



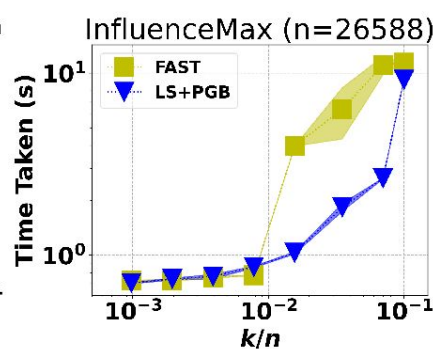




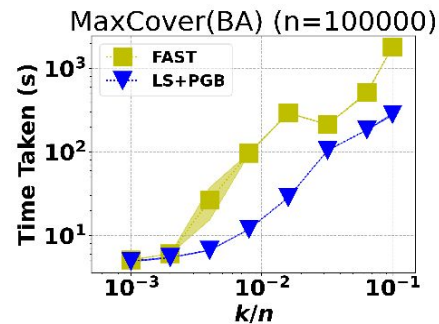
- Both **FAST** and **LS+PGB** exhibit a linear scaling behavior with the increasing k values
- Overall on an average **LS+PGB** achieves the objective in **less than half the total queries required by FAST** for all of the applications but **TrafficMonitor** and **InfluenceMax**.
- For **TrafficMonitor** & **InfluenceMax**, **FAST** requires **1.5** and **1.9** times the queries needed by **LS+PGB**

Empirical Results - Adaptive Rounds



- **LS+PGB** exhibits a very good scaling behavior with the increasing k values with at most 5 fold increase in adaptive rounds with 100 fold increase in k value.
- Overall on an average **FAST** requires more than 3.5 times the adaptive rounds needed by **LS+PGB** to achieve the objective.
- For **MaxCover**, **RevenueMax**, **TwitterSumm** and **ImageSumm** FAST requires 3.5, 4.3, 4.8 and 7.2 times more adaptive rounds.

Empirical Results - Time Taken



- Both algorithms exhibit **linear scaling of runtime with k**
- On many instances, **LS+PGB is faster by more than an order of magnitude**
- Overall on an average **FAST** requires almost **4.8 times the time needed by LS+PGB** to achieve the objective.
- For **TwitterSumm**, **MaxCover**, **RevenueMax** and **ImageSumm** **FAST** is on average **4.6**, **4.7**, **6.8** and **19.2** times slower than **LS+PGB**.

Empirical Results - Overall Result

Application	Runtime (s)		Objective Value		Queries	
	FAST	LS+PGB	FAST	LS+PGB	FAST	LS+PGB
TrafficMonitor	3.7×10^{-1}	2.1×10^{-1}	4.7×10^8	5.0×10^8	3.5×10^3	2.4×10^3
InfluenceMax	4.4×10^0	2.3×10^0	1.1×10^3	1.1×10^3	7.7×10^4	4.0×10^4
TwitterSumm	1.6×10^1	3.5×10^0	3.8×10^5	3.8×10^5	1.5×10^5	6.2×10^4
RevenueMax	3.9×10^2	5.4×10^1	1.4×10^4	1.4×10^4	7.6×10^4	2.7×10^4
MaxCover (BA)	3.7×10^2	7.6×10^1	6.0×10^4	6.3×10^4	5.8×10^5	1.8×10^5
ImageSumm	1.6×10^1	8.1×10^{-1}	9.1×10^3	9.1×10^3	1.3×10^5	4.8×10^4

Conclusion

- In this paper, we made the following contributions:
 - Theoretical
 - **LinearSeq**: A constant-factor algorithm for SM with smaller adaptivity than any previous algorithm, especially for values of k that are large relative to n .
 - **ThresholdSeq**: An algorithm that adds elements that have a gain of a specified threshold with expected linear query complexity and logarithmic adaptive rounds.
 - **LS+PGB**: A parallelized greedy algorithm which is used in conjunction with LinearSeq and ThresholdSeq. It obtains nearly the optimal result, in terms of ratio, adaptivity and query complexity.
 - Empirical
 - LS+PGB is faster than the state-of-art algorithm FAST in an extensive empirical evaluation.