
Fast Feature Selection with Fairness Constraints

Francesco Quinzan¹ Rajiv Khanna² Moshik Hershcovitch³ Sarel Cohen⁴ Daniel Waddington³
Tobias Friedrich⁴ Michael W. Mahoney^{5,6}

Abstract

We study the fundamental problem of selecting *optimal* features for model construction. This problem is computationally challenging on large datasets, even with the use of greedy algorithm variants. To address this challenge, we extend the adaptive query model, recently proposed for the greedy forward selection for submodular functions, to the faster paradigm of Orthogonal Matching Pursuit for non-submodular functions. The proposed algorithm achieves exponentially fast parallel run time in the adaptive query model, scaling much better than prior work. Furthermore, our extension allows the use of downward-closed constraints, which can be used to encode certain fairness criteria into the feature selection process. We prove strong approximation guarantees for the algorithm based on standard assumptions. These guarantees are applicable to many parametric models, including Generalized Linear Models. Finally, we demonstrate empirically that the proposed algorithm competes favorably with state-of-the-art techniques for feature selection, on real-world and synthetic datasets.

1. Introduction

We study the fundamental problem of selecting a few features out of many for a given modeling problem, while satisfying additional side constraints. This setup can be used to encode certain notions of fairness in feature selection in a principled way.¹ Formally, given a function $l : \mathbb{R}^n \rightarrow \mathbb{R}_{>0}$ expressing the goodness of fit, we search for a set of features

¹Department of Computer Science, University of Oxford ²Purdue University ³IBM ⁴Hasso Plattner Institute ⁵ICSI ⁶UC Berkeley. Correspondence to: Francesco Quinzan <francesco.quinzan@cs.ox.ac.uk>.

Presented at the 2nd Workshop on Formal Verification of Machine Learning, co-located with the 40th International Conference on Machine Learning, Honolulu, Hawaii, USA., 2023. Copyright 2023 by the author(s).

¹In this paper, we consider several fairness constraints proposed in the literature, but there may be other notions of fairness which do not fall within our theoretical framework.

S maximizing the function

$$f(S) := l(\beta^{(S)}) - l(\mathbf{0}). \quad (1)$$

Here, $\mathbf{0}$ represents the n -dimensional zero vector, and $\beta^{(S)}$ is a vector maximizing $l(\cdot)$ with support in S . If we denote by \mathcal{I} the set of all solution sets that satisfy the side constraints, then the feature selection optimization problem with side constraints can be formalized as

$$\arg \max_{S \subseteq [n]: S \in \mathcal{I}} f(S) = \arg \max_{S \subseteq [n]: S \in \mathcal{I}} l(\beta^{(S)}) - l(\mathbf{0}), \quad (2)$$

where $[n]$ is the index set of all the features. Often, \mathcal{I} corresponds to an r -sparsity constraint, i.e., a solution S is feasible if it contains at most r features.

Several algorithms have been proposed for feature selection under sparsity constraints. Some examples include forward step-wise selection methods (Elenberg et al., 2018; Qian & Singer, 2019), the Orthogonal Matching Pursuit (Needell & Tropp, 2010; Elenberg et al., 2018; Sakaue, 2020), forward-backward methods (Jalali et al., 2011; Liu et al., 2014), Pareto optimization (Qian et al., 2015; 2020), Exponential Screening (Rigollet & Tsybakov, 2010), and gradient-based methods (Jain et al., 2014; Yuan et al., 2017). These algorithms, however, are computationally inefficient on large datasets. Furthermore, there are a limited number of studies that take into account additional side constraints (e.g., matroids (Chen et al., 2018; Gatmiry & Gomez-Rodriguez, 2018)), which can be crucial when deploying machine learning systems in the real world.

Recently, there has been a growing effort towards developing fair algorithms for many fundamental problems, such as regression and classification (Zafar et al., 2017a; Kim et al., 2018; Feldman et al., 2015; Agarwal et al., 2018; Grgic-Hlaca et al., 2018b), matching (Chierichetti et al., 2019), and summarization (Halabi et al., 2020). Several definitions of fairness can be incorporated in the learning process as additional side constraints (Grgic-Hlaca et al., 2018b;a; Grgic-Hlaca et al., 2016; Chierichetti et al., 2019; Halabi et al., 2020; Zafar et al., 2017a; Woodworth et al., 2017; Donini et al., 2018; Agarwal et al., 2018; 2019). Motivated by this line of research, we consider the following question: *How can we efficiently perform feature selection, while taking into account additional constraints such as fairness?*

We address this question by proposing a novel algorithm that combines the paradigms of matching pursuit together with the adaptive query model for the problem (2). This algorithm is much faster than previously known techniques. At the same time, it allows to incorporate certain notions of fairness in the learning process, via a reduction to the p -system side constraints (see Jenkyns (1976) and Section 2.2).

Our Contribution. We propose a novel matching pursuit algorithm for the constrained feature selection problem (2). This algorithm uses oracle access to the gradient $\nabla l(\beta^{(S)})$, and to an oracle for the evaluation of the feasibility of an input solution set. The feasibility oracle is well-aligned with previous related work (Elenberg et al., 2018; Sakaue, 2020). The main technical contributions of this paper are two-fold: (i) we extend the adaptive query model to a class of non-submodular objectives using the paradigm of gradient based pursuit algorithms; and (ii) we incorporate general downward-closed constraints in the optimization process beyond the standard sparsity constraints.

2. Preliminaries

Notation. We denote with n the number of features, i.e., the dimension of the domain of $l(\cdot)$, and we define $[n] := \{1, 2, \dots, n\}$. For any $s \in [n]$, we denote with \mathbf{e}_s the unit vector, with a 1 for the coefficient indexed by s , and 0 otherwise. Feature sets are represented by sans script fonts, i.e., S, T . Vectors are represented by lower-case bold letters as $\mathbf{x}, \mathbf{y}, \beta$ and matrices are represented by upper-case bold letters, i.e., $\mathbf{X}, \mathbf{Y}, \Sigma$. For a feature set S , we denote with $\beta^{(S)}$ a vector maximizing $l(\cdot)$ with non-zero entries indexed by the set S . For a feature set T and a parameter vector β , we define $\nabla l(\beta)_T := \langle \nabla l(\beta), \sum_{s \in T} \mathbf{e}_s \rangle$. We denote with OPT the optimal value attained by the function $f(\cdot)$ as in (2). We denote with \mathcal{I} the p -system side constraint, and with r its rank, as defined in Section 2.1. The notation $\text{Cond}(T)$ denotes the set $\{s \in [n] \setminus T : T \cup \{s\} \in \mathcal{I}\}$.

2.1. Problem Formulation

We study optimization tasks as in the problem (2) under some additional assumptions on $l(\cdot)$, which are often satisfied in practical applications. We specifically consider functions that are (M, m) -(smooth, restricted concave) of a function (Negahban et al., 2010) (see Appendix A). Furthermore, we model the side constraints as p -systems (see Appendix B).

2.2. Embedding Certified Notions of Fairness via p -Systems

Procedural fairness metrics. Procedural fairness focuses on selecting features based on perceived fairness as envisioned by human beings during the decision-making process, rather than on the fairness of the outcome. This notion of

fairness is measured by gauging “the degree to which people consider various features to be fair” (Grgic-Hlaca et al., 2016). This is in contrast to measuring fairness of the *outcomes* of such decisions, for example, by down weighing decisions that affect users of protected groups (e.g., race or gender).

In this work, we consider measures for procedural fairness studied by (Grgic-Hlaca et al., 2016) and (Grgic-Hlaca et al., 2018b). However, our framework is not specific to these definitions. These measures consist of monotone set functions $h : 2^{[n]} \rightarrow [0, 1]$. For an input feature set $T \subseteq [n]$, the value $h(T)$ quantifies the perceived fairness of T , with $h(T) = 0$ corresponding to maximum fairness and $h(T) = 1$ corresponding to maximum unfairness. We refer the reader to (Grgic-Hlaca et al., 2016; Grgic-Hlaca et al., 2018b) for a precised definition of these concepts. While procedural fairness may not imply fairness of the outcome, it has been observed that in some cases procedurally fair feature sets maintain good outcome fairness (Grgic-Hlaca et al., 2016; Grgic-Hlaca et al., 2018b).

Feature partitions. Our proposed approach also includes as a special case the framework proposed by (Celis et al., 2018). Here, features are grouped into disjoint clusters $[n] = X_1 \cup \dots \cup X_\ell$. The constraints are specified using λ_j which encodes the maximum number of features that can be selected from cluster X_j . In other words, feature sets $T \subseteq [n]$ is feasible if the number of data-points intersecting a class X_j does not exceed the corresponding threshold λ_j . Formally, we define the set of constraints $\mathcal{I}_{\text{cl}}^\lambda := \{T \subseteq [n] : |T \cap X_j| \leq \lambda_j \forall j \in [\ell]\}$. This set of constraints is a *matroid*, which is a p -system with $p = 1$.

A generalization of the aforementioned partition matroid was considered by (Halabi et al., 2020). For each element in any partition set X_j , we are given a lower- and an upper-bound on the number of elements that can be selected from this set. Bounds are denoted by ℓ_j and u_j respectively. The set of constraints can be written as $\mathcal{I}_{\text{p}} := \{T \subseteq [n] : \ell_j \leq |T \cap X_j| \leq c_j \forall j \in [\ell]\}$. This set of constraints is, in general, not a p -system. However, (Halabi et al., 2020) show that one can consider a relaxation of the constraint set $\bar{\mathcal{I}}_{\text{p}} := \{T \subseteq [n] : T \subseteq S \text{ for any set } S \in \mathcal{I}_{\text{p}}\}$, which is equivalent from the optimization perspective. Any monotone optimization objective (as in (1)) yields the same solution sets on \mathcal{I}_{p} and $\bar{\mathcal{I}}_{\text{p}}$. The set of constraints $\bar{\mathcal{I}}_{\text{p}}$ is a *matroid*, i.e., a p -system with $p = 1$ (Edmonds, 1970).

3. Algorithmic Overview

Our algorithm, which we call FAST_{OMP} , is presented in Algorithm 1. This algorithm is based on a technique called adaptive sequencing (Balkanski et al., 2019; Breuer et al., 2020), which was recently proposed for highly scalable maximization of submodular functions.

Say X is the complete set of candidate features, and

Algorithm 1 FASTOMP

$S \leftarrow \emptyset$;
while the number of iterations is less than ε^{-1} and $\text{Cond}(S) \neq \emptyset$ **do**
 $X \leftarrow \{s \in [n]: \{s\} \cup S \in \mathcal{I}\}$;
 $t \leftarrow (1 - \varepsilon) \frac{m}{|\mathcal{T}|M} \|\nabla l(\beta^{(S)})_{\mathcal{T}}\|_2^2$ with $\mathcal{T} \subseteq X$ maximizing $\|\nabla l(\beta^{(S)})_{\mathcal{T}}\|_2^2$ s.t. $|\mathcal{T}| \leq r$;
while $X \neq \emptyset$ and $\text{Cond}(S) \neq \emptyset$ **do**
 $\{a_1, a_2, \dots, a_k\} \leftarrow \text{RNDSEQ}(X, S)$ and define $S_j \leftarrow S \cup \{a_1, \dots, a_j\}$;
observe $X_j \leftarrow \{s \in X: \langle \nabla l(\beta^{(S_j)}), \mathbf{e}_s \rangle^2 \geq t \text{ and } s \in \text{Cond}(S_j)\}$;
 $j^* \leftarrow \min_j \{j: |X_j| < (1 - \varepsilon)|X|\}$
 $X \leftarrow X_{j^*}$ and $S \leftarrow S_{j^*}$;
end while
end while **return** S ;

S is the current solution. Starting from $S \leftarrow \emptyset$, the FASTOMP iteratively generates a random sequence of features $\{a_1, a_2, \dots, a_k\}$ with the RNDSEQ sub-routine (details in Appendix H), such that the set $\{a_1, a_2, \dots, a_k\} \cup S$ is a maximal independent set of \mathcal{I} . After a sequence is generated, the FASTOMP identifies a prefix $\{a_1, \dots, a_{j^*}\}$ that is added to the current solution. The index j^* defining this prefix is chosen such that it holds $|X_j| \geq (1 - \varepsilon)|X|$, for all $0 \leq j < j^*$. This inequality ensures that any point added to the current solution yields $\langle \nabla l(\beta^{(S)}), \mathbf{e}_s \rangle^2 \geq t$ in expected value. Finally, the ground set X is updated as to include only those points that yield a good improvement to the new solution. The RNDSEQ sub-routine used to generate $\{a_1, a_2, \dots, a_k\}$ corresponds to Algorithm A by Karp et al. (1988). Here, k is the size of the independent set returned in the current iteration.

4. Approximation Guarantees

In this section, we study the approximation guarantees for Algorithm 1, when solving Problem (2).

Theorem 4.1. *Define the support selection function $f(\cdot)$ as in (1), for the given function $l(\cdot)$ that is (M, m) -restricted smooth, restricted strong concave, on the sparse sub-domain Ω_{2r} . Consider a p -system \mathcal{I} of rank r over $[n]$, and let S^* be the output of Algorithm 1 while OPT is the optimum solution set for the Problem 2. Then,*

$$\frac{\mathbb{E}[f(S^*)]}{\text{OPT}} \geq \frac{1}{1+p} \left(1 - \exp \left\{ -(1-\varepsilon)^2 \frac{m^3}{M^3} \right\} \right),$$

for all $0 < \varepsilon < 1$. Furthermore, in the specific case when \mathcal{I} is r -sparsity constraint over $[n]$, then,

$$\frac{\mathbb{E}[f(S^*)]}{\text{OPT}} \geq \left(1 - \exp \left\{ -(1-\varepsilon)^2 \frac{m^2}{M^2} \right\} \right).$$

A full proof of this theorem is deferred to Appendix J. The parameter p in Theorem 4.1 is always upper-bounded by the maximum number of features r that we wish to select for model construction.

We also provide bounds for the run time of the FASTOMP as follows.

Theorem 4.2. *Algorithm 1 terminates after $\mathcal{O}(\varepsilon^{-2} \log n)$ rounds of calls to the oracle function, and it uses at most $\mathcal{O}(\varepsilon^{-2} r \log n)$ oracle queries. Furthermore, Algorithm 1 requires expected $\mathcal{O}(\varepsilon^{-2} \sqrt{r} \log n)$ independent calls to the oracle for the p -system \mathcal{I} , and the total expected number of calls to the oracle for the p -system \mathcal{I} is $\mathcal{O}(\varepsilon^{-2} nr \log n)$.*

The proof of Theorem 4.2 is deferred to Appendix K. The estimates on the rounds of adaptivity extends to the PRAM model. If we denote with d_l the depth required to evaluate the oracle function on a set, then the FASTOMP has $\mathcal{O}(\varepsilon^{-2} d_l \log n)$ depth.

5. Experiments

5.1. Benchmarks

We compare our algorithm against the Orthogonal Matching Pursuit (SDSOMP, Krause & Cevher (2010); Elenberg et al. (2018)), a submodular optimization algorithm (ISK, Iyer & Bilmes (2013)), the simple lasso (Lasso), and uniform random sampling (Random, see H).

5.2. Datasets

Synthetic Datasets. We generate a synthetic dataset with 10^6 features, 10^4 observations, and size ~ 19 GB. Features are sampled with a joint Gaussian distribution with mean vector $\mu = \mathbf{0}$, and a covariance matrix Σ designed to ensure that 10% of the features are highly correlated with the response, while the remaining features have low correlation with the response variable. We then add posterior uniform noise, and normalize the observations.

The ProPublica COMPAS DataSet. We consider the well-known ProPublica COMPAS dataset, which is a pretrial risk assessment instrument (Larson et al., 2016). This dataset was constructed in 2016, using data of defendants from Broward County, FL (US), who had been arrested in 2013 or 2014. This dataset was intended to be used to predict if a criminal was likely to re-offend, based on previous arrest charges and demographic information. Predictions based on the COMPAS datasets were found to be racially biased (Angwin et al., 2016; Berk et al., 2021) (see Appendix N for details).

5.3. Results

Synthetic Dataset. We search for a set of features maximizing the R^2 objective. We consider a randomly generated

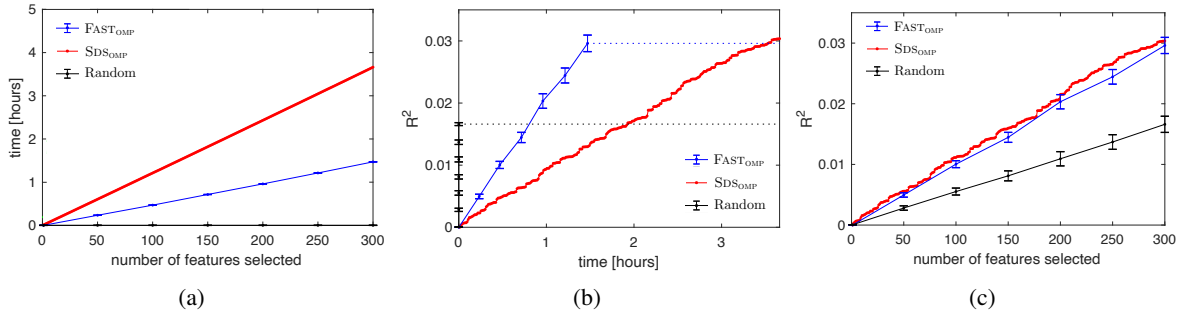


Figure 1. Results on the Synthetic Dataset (mid row), as described in Section 5. For the randomized algorithms, each dot corresponds to the sample mean of multiple runs, and error bars correspond to the standard deviation. Across datasets, we see that SDS_{OMP} achieves substantial speedup, while remaining competitive in terms of precision quality of features being selected.

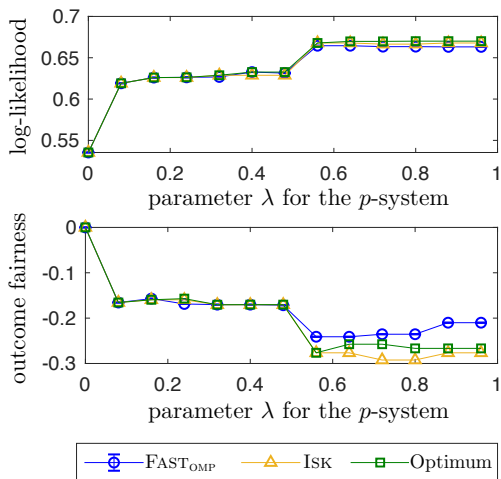


Figure 2. Experiments on the COMPAS Dataset, as detailed in Section 5. For the FAST_{OMP} we perform multiple runs and we report on the sample mean. On this dataset, the variance in accuracy and outcome fairness is negligible for FAST_{OMP}.

p -system on top of the features. To illustrate scalability on such a large dataset, we compare speed of feature selection while satisfying the side constraints (Figure 1(a)) and the corresponding solution accuracy with respect to time spent (Figure 1(b)) for selecting upto $k = 300$ features. We observe that our algorithm FAST_{OMP} is significantly faster while maintaining a good quality solution when compared to the SDS_{OMP} and the Random. In Figure 1(c), we further observe that the FAST_{OMP} and the SDS_{OMP} achieve similar solution quality for a given number of features.

ProPublica COMPAS Dataset. In this section, we reproduce the experiments by Grgic-Hlaca et al. (2018b) on the COMPAS dataset. We use regularized logistic regression to predict the recidivism risk and use fairness constraints given by the feature-apriori accuracy (Grgic-Hlaca et al., 2018b). These constraints are encoded as p -system side constraints $\mathcal{I}_{acc}^\lambda$ as described in Section

2.2. We report on the outcome fairness of each output feature set, by estimating the racial bias of the corresponding classifier. Following Grgic-Hlaca et al. (2018b); Kleinberg et al. (2017); Zafar et al. (2017a), we examine the false positive (FPR) and false negative (FNR) rates for whites (w) and non-whites (nw) as *outcome fairness* = $-|FPR_w - FPR_{nw}| - |FNR_w - FNR_{nw}|$. This measure of fairness varies between -2 and 0 , with -2 corresponding to maximum unfairness and 0 to maximum fairness.

The results for this set of experiments are displayed in Figure 2, where we plot the accuracy and outcome fairness as a function of the parameter λ for the constraints $\mathcal{I}_{acc}^\lambda$. We compare the solution quality and fairness of the FAST_{OMP}, against the ISK, and the solution with best possible accuracy (optimum). We observe that the FAST_{OMP} achieves nearly optimal solution. Although the outcome fairness inevitably decreases for increasing process unfairness, the FAST_{OMP} maintains a more graceful degradation than the other algorithms.

6. Ethics Discussion

A major hurdle in further research into fair learning is the lack of gold standard benchmarks. The COMPAS dataset is used for empirical evaluations in several studies. However, its use for benchmarking has also been criticized (Bao et al., 2021b). Furthermore, over-tuning notions of fairness for a single dataset could be problematic.

Fairness criteria should always take into account the contextual grounding of the dataset, and the trained model that operates within. As such, the mathematical formalism of fairness as constraints may evolve. We hope that our framework motivates further research into addressing the above concerns about fairness in an algorithmically scalable way.

Acknowledgments. This project received partial support from ELSA: European Lighthouse on Secure and Safe AI project (agreement No. 101070617 under UK guarantee).

References

- Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., and Wallach, H. M. A reductions approach to fair classification. In *Proc. of ICML*, volume 80, pp. 60–69, 2018.
- Agarwal, A., Dudík, M., and Wu, Z. S. Fair regression: Quantitative definitions and reduction-based algorithms. In *Proc. of ICML*, volume 97, pp. 120–129, 2019.
- Angwin, J., Larson, J., Mattu, S. a., and Kirchner, L. There’s software used across the country to predict future criminals. And it’s biased against blacks, 2016.
- Bahmani, S., Raj, B., and Boufounos, P. T. Greedy sparsity-constrained optimization. *Journal of Machine Learning Research*, 14(1):807–841, 2013.
- Balkanski, E. and Singer, Y. The adaptive complexity of maximizing a submodular function. In *Proc. of STOC*, pp. 1138–1151, 2018.
- Balkanski, E., Rubinstein, A., and Singer, Y. An optimal approximation for submodular maximization under a matroid constraint in the adaptive complexity model. In *Proc. of STOC*, pp. 66–77, 2019.
- Bamler, R., Zhang, C., Opper, M., and Mandt, S. Perturbative black box variational inference. In *Proc. of NPIS*, pp. 5079–5088, 2017.
- Bandeira, A. S., Dobriban, E., Mixon, D. G., and Sawin, W. F. Certifying the restricted isometry property is hard. *IEEE Transactions of Information Theory*, 59(6):3448–3450, 2013.
- Bao, M., Zhou, A., Zottola, S., Brubach, B., Desmarais, S., Horowitz, A., Lum, K., and Venkatasubramanian, S. It’s compaslicated: The messy relationship between RAI datasets and algorithmic fairness benchmarks. *CoRR*, abs/2106.05498, 2021a.
- Bao, M., Zhou, A., Zottola, S. A., Brubach, B., Desmarais, S., Horowitz, A. S., Lum, K., and Venkatasubramanian, S. It’s COMPASlicated: The messy relationship between RAI datasets and algorithmic fairness benchmarks. In *Proc. of NeurIPS Datasets and Benchmarks Track*, 2021b.
- Berk, R., Heidari, H., Jabbari, S., Kearns, M., and Roth, A. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50(1):3–44, 2021.
- Borodin, A., von zur Gathen, J., and Hopcroft, J. E. Fast parallel matrix and GCD computations. *Information and Control*, 52(3):241–256, 1982.
- Breuer, A., Balkanski, E., and Singer, Y. The FAST algorithm for submodular maximization. In *Proc. of ICML*, pp. 1134–1143, 2020.
- Celis, L. E., Keswani, V., Straszak, D., Deshpande, A., Kathuria, T., and Vishnoi, N. K. Fair and diverse dpp-based data summarization. In *Proc. of ICML*, pp. 715–724, 2018.
- Chekuri, C. and Quanrud, K. Parallelizing greedy for submodular set function maximization in matroids and beyond. In *Proc. of STOC*, pp. 78–89, 2019.
- Chen, L., Feldman, M., and Karbasi, A. Weakly submodular maximization beyond cardinality constraints: Does randomization help greedy? In *Proc. of ICML*, pp. 804–813, 2018.
- Chierichetti, F., Kumar, R., Lattanzi, S., and Vassilvitskii, S. Matroids, matchings, and fairness. In *Proc. of AISTATS*, pp. 2212–2220, 2019.
- Chistov, A. L. Fast parallel calculation of the rank of matrices over a field of arbitrary characteristic. In *Proc. of FCT*, pp. 63–69, 1985.
- Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163, 2017.
- Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. Algorithmic decision making and the cost of fairness. In *Proc. of KDD*, pp. 797–806, 2017.
- Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. Algorithmic decision making and the cost of fairness. In *Proc. of KDD*, pp. 797–806, 2017.
- Cowgill, B. and Tucker, C. E. Economics, fairness and algorithmic bias. *preparation for: Journal of Economic Perspectives*, 2019.
- Das, A. and Kempe, D. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In *Proc. of ICML*, pp. 1057–1064, 2011.
- Ding, C. H. Q. and Peng, H. Minimum redundancy feature selection from microarray gene expression data. *Journal of Bioinformatics and Computational Biology*, 3(2):185–206, 2005.
- Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J., and Pontil, M. Empirical risk minimization under fairness constraints. In *Proc. of NeurIPS*, pp. 2796–2806, 2018.
- Edmonds, J. *Submodular functions, matroids, and certain polyhedra.*, pp. 69–87. 1970.
- Elenberg, E. R., Khanna, R., Dimakis, A. G., and Negahban, S. N. Restricted strong convexity implies weak submodularity. *The Annals of Statistics*, 46(6B):3539–3568, 2018.

- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. Certifying and removing disparate impact. In *Proc. of SIGKDD*, pp. 259–268, 2015.
- Gatmiry, K. and Gomez-Rodriguez, M. On the network visibility problem. *CoRR*, abs/1811.07863, 2018.
- Grgic-Hlaca, N., Zafar, M. B., Gummadi, K., and Weller, A. The case for process fairness in learning: Feature selection for fair decision making. In *NeurIPS Symposium on Machine Learning and the Law*, 2016.
- Grgic-Hlaca, N., Redmiles, E. M., Gummadi, K. P., and Weller, A. Human perceptions of fairness in algorithmic decision making: A case study of criminal risk prediction. In *Proc. of WWW*, pp. 903–912, 2018a.
- Grgic-Hlaca, N., Zafar, M. B., Gummadi, K. P., and Weller, A. Beyond distributive fairness in algorithmic decision making: Feature selection for procedurally fair learning. In *Proc. of AAAI*, pp. 51–60, 2018b.
- Halabi, M. E., Mitrovic, S., Norouzi-Fard, A., Tardos, J., and Tarnawski, J. Fairness in streaming submodular maximization: Algorithms and hardness. In *Proc. of NeurIPS*, 2020.
- Ibarra, O. H., Moran, S., and Rosier, L. E. A note on the parallel complexity of computing the rank of order n matrices. *Information Processing Letters*, 11(4/5):162, 1980.
- Iyer, R. K. and Bilmes, J. A. Submodular optimization with submodular cover and submodular knapsack constraints. In *Proc. of NIPS*, pp. 2436–2444, 2013.
- Jain, P., Tewari, A., and Kar, P. On iterative hard thresholding methods for high-dimensional m -estimation. In *Proc. of NIPS*, pp. 685–693, 2014.
- Jalali, A., Johnson, C. C., and Ravikumar, P. On learning discrete graphical models using greedy methods. In *Proc. of NIPS*, pp. 1935–1943, 2011.
- Jenkyns, T. A. The efficacy of the "greedy" algorithm. In *Proc. of SEICCGTC*, pp. 341–350, 1976.
- Kalimeris, D., Kaplun, G., and Singer, Y. Robust influence maximization for hyperparametric models. In *Proc. of ICML*, pp. 3192–3200, 2019.
- Karp, R. M., Upfal, E., and Wigderson, A. The complexity of parallel search. *Journal of Computer and System Science*, 36(2):225–253, 1988.
- Kim, M. P., Reingold, O., and Rothblum, G. N. Fairness through computationally-bounded awareness. In *Proc. of NeurIPS*, pp. 4847–4857, 2018.
- Kleinberg, J. M., Mullainathan, S., and Raghavan, M. Inherent trade-offs in the fair determination of risk scores. In *Proc. of ITCS*, volume 67, pp. 43:1–43:23, 2017.
- Krause, A. and Cevher, V. Submodular dictionary selection for sparse representation. In *Proc. of ICML*, pp. 567–574, 2010.
- Larson, J., Roswell, M., and Atlidakis, V. Data and analysis for ‘how we analyzed the COMPAS recidivism algorithm’. <https://github.com/propublica/compas-analysis>, 2016.
- Liu, J., Ye, J., and Fujimaki, R. Forward-backward greedy algorithms for general convex smooth functions over A cardinality constraint. In *Proc. of ICML*, volume 32, pp. 503–511, 2014.
- Mairal, J. and Yu, B. Complexity analysis of the lasso regularization path. In *Proc. of ICML*, 2012.
- Mirzasoleiman, B., Badanidiyuru, A., and Karbasi, A. Fast constrained submodular maximization: Personalized data summarization. In *Proc. of ICML*, volume 48, pp. 1358–1367, 2016.
- Mulmuley, K. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. *Combinatorica*, 7(1): 101–104, 1987.
- Needell, D. and Tropp, J. A. Cosamp: iterative signal recovery from incomplete and inaccurate samples. *Communication of the ACM*, 53(12):93–100, 2010.
- Negahban, S., Ravikumar, P., Wainwright, M., and Yu, B. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Statistical Science*, 27:538–557, 10 2010.
- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. An analysis of approximations for maximizing submodular set functions - II. *Mathematical Programming*, 14(3): 73–87, 1978.
- Nowozin, S. Debiasing evidence approximations: On importance-weighted autoencoders and jackknife variational inference. In *Proc. of ICLR*, 2018.
- Qian, C., Yu, Y., and Zhou, Z. Subset selection by pareto optimization. In *Proc. of NIPS*, pp. 1774–1782, 2015.
- Qian, C., Bian, C., and Feng, C. Subset selection by pareto optimization with recombination. In *Proc. of AAAI*, pp. 2408–2415, 2020.
- Qian, S. and Singer, Y. Fast parallel algorithms for statistical subset selection problems. In *Proc. of NeurIPS*, pp. 5073–5082, 2019.

- Quinzan, F., Doskoc, V., Göbel, A., and Friedrich, T. Adaptive sampling for fast constrained maximization of sub-modular functions. In *Proc. of AISTATS*, pp. 964–972, 2021.
- Rigollet, P. and Tsybakov, A. Exponential screening and optimal rates of sparse estimation. *The Annals of Statistics*, 39:731–771, 2010.
- Rudin, C., Wang, C., and Coker, B. The age of secrecy and unfairness in recidivism prediction. *arXiv:1811.00731*, 2018.
- Sakaue, S. On maximization of weakly modular functions: Guarantees of multi-stage algorithms, tractability, and hardness. In *Proc. of AISTATS*, pp. 22–33, 2020.
- Torkkola, K. Feature extraction by non-parametric mutual information maximization. *Journal of Machine Learning Research*, 3:1415–1438, 2003.
- Waddington, D., Dickey, C., Hershcovitch, M., and Seshadri, S. An architecture for memory centric active storage (MCAS). *CoRR*, abs/2103.00007, 2021a.
- Waddington, D., Dickey, C., Xu, L., Hershcovitch, M., and Seshadri, S. A high-performance persistent memory key-value store with near-memory compute. *CoRR*, abs/2104.06225, 2021b.
- Woodworth, B. E., Gunasekar, S., Ohannessian, M. I., and Srebro, N. Learning non-discriminatory predictors. In *Proc. of COLT*, volume 65, pp. 1920–1953, 2017.
- Yuan, X., Li, P., and Zhang, T. Gradient hard thresholding pursuit. *Journal of Machine Learning Research*, 18:166:1–166:43, 2017.
- Zafar, M. B., Valera, I., Gomez-Rodriguez, M., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *Proc. of AISTATS*, volume 54, pp. 962–970, 2017a.
- Zafar, M. B., Valera, I., Gomez-Rodriguez, M., and Gummadi, K. P. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proc. of WWW*, pp. 1171–1180, 2017b.
- Zafar, M. B., Valera, I., Gomez-Rodriguez, M., Gummadi, K. P., and Weller, A. From parity to preference-based notions of fairness in classification. In *Proc. of NIPS*, pp. 229–239, 2017c.
- Zhao, Z., Anand, R., and Wang, M. Maximum relevance and minimum redundancy feature selection methods for a marketing machine learning platform. In *Proc. of DSAA*, pp. 442–452, 2019.

A. Definition of Restricted Strong Concavity and Restricted Smoothness

We study optimization tasks as in the problem (2) under some additional assumptions on $l(\cdot)$, which are often satisfied in practical applications (see Appendix E). Define an r -sparse subdomain as a set of the form $\Omega_r := \{(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^n \times \mathbb{R}^n : \|\mathbf{x}\|_0 \leq r, \|\mathbf{y}\|_0 \leq r, \|\mathbf{x} - \mathbf{y}\|_0 \leq r\}$. We now define the notions of Restricted Strong Concavity (RSC) and Restricted Smoothness (RSM) of a function.

Definition A.1 (RSC, RSM (Negahban et al., 2010)). A function $l(\cdot)$ is said to be restricted strong concave (RSC) with parameter m and restricted smooth (RSM) with parameter M on a subdomain Ω_r iff, for all $(\mathbf{x}, \mathbf{y}) \in \Omega_r$ it holds that $-\frac{m}{2}\|\mathbf{y} - \mathbf{x}\|_2 \geq l(\mathbf{y}) - l(\mathbf{x}) - \langle \nabla l(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \geq -\frac{M}{2}\|\mathbf{y} - \mathbf{x}\|_2$.

We say that l is (M, m) -(smooth, restricted concave), if it fulfills the conditions as in Definition A.1 with parameters M and m . RSC/RSM often hold in practice, we refer the reader to Appendix E for further discussion of these properties.

B. Definition of p -System

In order to give an axiomatic definition of p -systems, we introduce additional terminology. Given a collection of feasible solutions \mathcal{I} over a ground set $[n]$ and a set $T \subseteq [n]$, we denote with $\mathcal{I}|_T$, the restricted feasible solution set, as the collection consisting of all sets $S \subseteq T$ s.t. $S \in \mathcal{I}$. We define $\text{Cond}(T)$ as the set of all $s \in [n] \setminus T$ such that $T \cup \{s\} \in \mathcal{I}$. A set T is a maximal independent set if it holds that $\text{Cond}(T) = \emptyset$. A base for \mathcal{I} is any maximal set $T \in \mathcal{I}$.

Definition B.1 (p -Systems (Jenkyns, 1976)). A p -system \mathcal{I} over $[n]$ is a collection of subsets of $[n]$ such that: (i) $\emptyset \in \mathcal{I}$; (ii) for any two sets $S \subseteq T \subseteq [n]$, if $T \in \mathcal{I}$ then $S \in \mathcal{I}$; (iii) for any set $T \subseteq [n]$ and any bases $S, U \in \mathcal{I}|_T$ it holds $|S| \leq p|U|$.

The second defining axiom is commonly referred to as subset-closure or downward-closed property. The rank r of a p -system \mathcal{I} is defined as the maximum cardinality of any feasible solution $T \in \mathcal{I}$.

C. Motivating Example

In this section, we motivate our analysis by showing that the adaptive sequencing prototype by (Balkanski et al., 2019) does not work for feature selection. this example can also be used to show that similar algorithms, such as FAST (Breuer et al., 2020) do not work for feature selection.

The adaptive sequencing prototype is presented in Algorithm 2. We refer to this algorithm as the ADAPTIVE SEQUENCING. Starting from the empty set, the ADAPTIVE SEQUENCING generates at every iteration a uniformly random sequence $\{a_1, \dots, a_k\}$ of the elements X not yet discarded. This sequence can be sampled uniformly at random, or using the RNDSEQ sub-routine described in Appendix H. Afterwards, the ADAPTIVE SEQUENCING filters elements that have a high marginal contribution, when added to S , and it determines the prefix of $\{a_1, \dots, a_k\}$ that is added to the current solution S . This prefix has the property that there is a large fraction of elements in X with high contribution to the current solution S .

Following and example provided by (Elenberg et al., 2018), we show that the ADAPTIVE SEQUENCING fails on a simple linear regression task with three features. For a fixed parameter $z > 0$, consider the following variables:

$$\begin{aligned} \mathbf{y} &= [1, 0, 0]^T \\ \mathbf{x}_1 &= [0, 1, 0]^T \\ \mathbf{x}_2 &= [z, \sqrt{1-z^2}, 0]^T \\ \mathbf{x}_3 &= [\delta z, 0, \sqrt{1-\delta^2 z^2}]^T \end{aligned}$$

Note that all variables have unit norm. Our goal is to choose two of the three variables $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ that best estimate \mathbf{y} , with respect to the R^2 objective. To this end, we introduce additional notation. For a given index set $S \subseteq [3]$, we denote with \mathbf{X}_S the matrix whose columns consists of the features indexed by S . For instance, for $S = \{1, 3\}$ it holds

$$\mathbf{X}_{\{1,3\}} = \begin{bmatrix} 0 & \delta z \\ 1 & 0 \\ 0 & \sqrt{1-\delta^2 z^2} \end{bmatrix}.$$

With this notation, we define the objective function for our problem as

$$f(S) = R^2(\beta^{(S)}) - R^2(\mathbf{0}) = (\mathbf{y}^T \mathbf{X}_S)(\mathbf{X}_S^T \mathbf{X}_S)^{-1}(\mathbf{X}_S^T \mathbf{y}), \quad (3)$$

Algorithm 2 ADAPTIVE SEQUENCING

```

S ← ∅, t ← maxe∈[n] f(e);
for Δ iterations do
  X ← [n];
  while X ≠ ∅ do
    {a1, a2, ..., ak} ← RNDSEQ(X, S);
    Xi ← {e ∈ X: fS ∪ {a1, ..., ai}(e) ≥ t and S ∪ {a1, ..., ai, e} ∈ ℐ};
    i* ← min{i: |Xi| ≤ (1 - ε)|X|} S ← S ∪ {a1, ..., ai};
    X = Xi*;
  end while
  t ← (1 - ε)t;
end for
return S;

```

with $R^2(\cdot)$ the R^2 objective evaluated on the model for an input parameter vector $\beta^{(S)}$. Using this formula, one can easily see that it holds

$$\begin{aligned}
 f(\{1\}) &= 0 & f(\{1, 2\}) &= 1 \\
 f(\{2\}) &= z^2 & f(\{1, 3\}) &= \delta^2 z^2 \\
 f(\{3\}) &= \delta^2 z^2 & f(\{2, 3\}) &= (1 + \delta)z^2 + \delta^2 z^4
 \end{aligned}$$

Clearly, the optimal solution is $f(\{1, 2\})$. However, on this instance the ADAPTIVE SEQUENCING outputs the solution $S = \emptyset$, attaining an f -value of $f(S) = 0$. The following lemma holds.

Lemma C.1. *Consider the ADAPTIVE SEQUENCING optimizing the function $f(S)$ as in (3), over sets $S \subseteq [3]$ of size at most $|S| \leq k$ with $k = 2$. Then, for any constant $\varepsilon > 0$, the ADAPTIVE SEQUENCING outputs either the solution $\{1, 3\}$ or the solution $\{2, 3\}$ with probability at least $1 - (2/3)^{\varepsilon^{-1} \log 1/\delta}$. In particular, the expected approximation guarantee of ADAPTIVE SEQUENCING converges to zero, for $\delta, z \rightarrow 0$.*

Proof. At the beginning of the optimization process, the constant t is set to $t = \delta^2 z^2$ and a sequence $\{a_1, a_2\}$. This sequence yield $\{a_1\} = \{3\}$ at least with probability $1/3$. If $\{a_1\} = \{3\}$ sequence is sampled, then the point $\{3\}$ is added to the current solution. Otherwise, the value t decreases of a multiplicative factor of $(1 - \varepsilon)$, and a new sequence is sampled. As long as $t > z^2$, the ADAPTIVE SEQUENCING can only output a solution that contains the point $\{3\}$. Hence, the solution $\{1, 2\}$ can only be sampled after t decreases to a value $t < z$, which requires at least $\varepsilon^{-1} \log 1/\delta$ iterations of the outer loop. Hence, the probability of sampling the solution $\{1, 2\}$ can be upper-bounded by the probability that no sequence with $\{a_1\} = \{3\}$ is sampled during the first $\varepsilon^{-1} \log 1/\delta$ iterations of the outer-loop. This probability can be estimated as $2/3^{\varepsilon^{-1} \log 1/\delta}$. Hence, the ADAPTIVE SEQUENCING outputs either the solution $\{1, 3\}$ or the solution $\{2, 3\}$ with probability at least $1 - (2/3)^{\varepsilon^{-1} \log 1/\delta}$. \square

We remark that in this example, the FAST_{OMP} yields a constant-factor approximation guarantee in expected value. In fact, we can show that our algorithm outputs the optimal solution in one iteration, at least with constant probability. We can bound the optimal parameters and in Line 4 of Algorithm 1 in a similar fashion as in Proposition E.5, and obtain that m and M are bounded as $m \geq 1 - \sqrt{1 - z^2}$ and $M \leq 1 + \sqrt{1 - z^2}$. It follows that the parameter in Line 4 of Algorithm 1 is upper-bounded as $t \leq (1 - \varepsilon)(1 - \sqrt{1 - z^2})/2 \leq z^2$, for δ sufficiently small. Suppose now that at the beginning of the iteration, a sequence $\{a_1, a_2\} = \{2, 1\}$. Since $t \leq z^2$, then the entire sequence is added to the current solution $\{2, 1\}$ and the algorithm outputs the optimum. Note that the desired sequence is sampled with probability $1/6$. It follows that the FAST_{OMP} outputs the optimum at least with constant probability. Hence, the FAST_{OMP} maintains a constant-factor approximation guarantee.

D. Weak Submodularity

Consider a function l that is restricted smooth and restricted strong concave. It is well known that the corresponding function f as in (1) has weak diminishing return properties. These diminishing returns property is called *weak submodularity*, and it was first introduced by (Das & Kempe, 2011) to study statistical subset selection problems. Functions that exhibit weak submodularity, i.e., weakly submodular functions, are defined in terms of the submodularity ratio, as follows.

Definition D.1 (Weak submodularity, Definition 2.3 by (Das & Kempe, 2011)). Consider a function $f : 2^{[n]} \rightarrow \mathbb{R}_{\geq 0}$. The submodularity ratio is defined as the largest scalar γ such that

$$\sum_{j \in S} (f(L \cup \{j\}) - f(L)) \geq \gamma (f(L \cup S) - f(L)),$$

for all sets $L, S \in [n]$ such that $L \cap S \neq \emptyset$. We say that f is γ -weakly submodular if its submodularity ratio is γ .

There is a well-known connection between weak submodularity and Problem 2. This connection was discovered by (Elenberg et al., 2018), and it can be formalized as follows.

Theorem D.2 (Theorem 1 by (Elenberg et al., 2018)). Define f as in (1), with a function l that is $(m_{|U|+k}, \tilde{M}_{|U|+k})$ -(strongly concave, smooth) on $\Omega_{|U|+k}$, and $\tilde{M}_{|U|+1}$ smooth on $\Omega_{|U|+1}$. Fix a set $U \in [n]$, and denote with $\gamma_{U,k}$ the largest scalar such that

$$\sum_{j \in S} (f(L \cup \{j\}) - f(L)) \geq \gamma_{U,k} (f(L \cup S) - f(L)),$$

for all sets $L, S \in [n]$ such that $L \cap S \neq \emptyset$, $L \subseteq U$, $|S| \leq k$. Then, the constant $\gamma_{U,k}$ is lower-bounded as

$$\gamma_{U,k} \geq \frac{m_{|U|+k}}{\tilde{M}_{|U|+1}} \geq \frac{m_{|U|+k}}{M_{|U|+k}}.$$

E. Restricted Strong Concavity and Smoothness for Feature Selection

Given a set of observations and a parametric family of distributions $\{p(\cdot; \beta) \mid \beta \in \Omega\}$ with $\Omega \subseteq \mathbb{R}^n$, we wish to identify a vector of parameters β maximizing the goodness of fit for these observation, according to a chosen measure l . For generalized linear models, common measures for feature selection are restricted strong concave and restricted smooth. We study the log-likelihood and the coefficient of determination, although analogous results hold for other similar statistics (Qian & Singer, 2019; Das & Kempe, 2011).

Maximizing the log-conditional. Assuming that the response follows a distribution in an exponential family, the log-conditional can be written as

$$\log p(\mathbf{y} \mid \mathbf{X}; \beta) = h^{-1}(\tau) - Z(\mathbf{X}, \beta) + g(\mathbf{y}, \tau) \quad (4)$$

with Z the log-partition function, and τ the dispersion parameter. The log-conditional is commonly used for learning. In the case of the simple linear model, it is possible to derive approximation guarantees in terms of the eigenvalues of \mathbf{X} . Additional assumptions on the random design of \mathbf{X} ensure that the log-conditional objective is restricted smooth and restricted strong concave. We refer the reader to Appendix E.1 for a discussion on these results.

Maximizing the R^2 objective. Consider a linear model with a normalized response variable \mathbf{y} . The R^2 objective is defined as

$$R_{\mathbf{X}, \mathbf{y}}^2(\beta) := 1 - \mathbb{E} [(\mathbf{y} - \langle \mathbf{X}, \beta \rangle)^2]. \quad (5)$$

This function is a popular measure for the goodness of fit. The function $R_{\mathbf{X}, \mathbf{y}}^2(\beta)$ is restricted strong concave and restricted smooth, with parameters depending on the properties of the matrix \mathbf{X} . A discussion on these results is deferred to Appendix E.2.

E.1. Restricted Strong Concavity and Smoothness of the Log-Conditional

In this section we discuss results concerning the (M, m) -(smoothness, strong concavity) of the log-conditional for generalized linear models. Consider a feature matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ and response variable \mathbf{y} . Assuming that the response follows a distribution in an exponential family, the log-conditional can be written as

$$\log p(\mathbf{y} \mid \mathbf{X}; \beta) = h^{-1}(\tau) - Z(\mathbf{X}, \beta) + g(\mathbf{y}, \tau) \quad (6)$$

with Z the log-partition function, and τ the dispersion parameter. We give approximation guarantees for the objective

$$l(\beta) = \log p(\mathbf{y} \mid \mathbf{X}; \beta) - \eta \|\beta\|_2^2, \quad (7)$$

for some parameter $\eta \geq 0$. We show that the function l is restricted smooth and restricted strong concave under various assumptions. We start discussing the simplest case of a logistic regression. We then study the general case as in (6), under the assumption that l has a non-zero regularization term. We then conclude with the general case, i.e., no assumptions on the regularization term.

Logistic regression. The log-likelihood function for the logistic regression is defined as

$$l(\beta^{(S)}) := \sum_i y_i \langle \mathbf{X}_S, \beta \rangle - \log \left(1 - e^{\langle \mathbf{X}_S, \beta \rangle} \right), \quad (8)$$

with \mathbf{X}_S the matrix of all features indexed by S and y_i the i -th observation, i.e., the i -th coefficient of the response \mathbf{y} . For this class of log-likelihood functions, the following result holds.

Proposition E.1. *The log-likelihood function l for the logistic regression as in (8), is (m, M) -(restricted smooth, restricted strong concave) with parameters*

$$m := \min_S \lambda_{\min}(\mathbf{X}_S^T \mathbf{X}_S) \text{ and } M := \max_S \lambda_{\max}(\mathbf{X}_S^T \mathbf{X}_S).$$

Log-conditional with non-zero regularization term. We now study log-conditional functions as in (6), under the assumption that the corresponding objective l has a regularization term with parameter $\eta > 0$. We introduce additional notation to this end. For any feature set $T \subseteq [n]$, denote with \mathcal{P}_T an operator that takes as input vectors $\mathbf{x} \in \mathbb{R}^n$, and it replaces all indices in $[n] \setminus T$ of \mathbf{x} by 0. For any vector $\mathbf{x} \in \mathbb{R}^p$, we define $\mathbf{x}_T := \mathcal{P}_T(\mathbf{x})$. We consider the following assumption on the distribution of the features.

Assumption E.2 (Assumption by (Bahmani et al., 2013)). For fixed constants $r, R > 0$, we make the following assumption on the feature matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$. The rows \mathbf{x} of \mathbf{X} are generated *i.i.d.*, such that the following additional conditions hold. For any set $T \subseteq [p]$ of size $|T| \leq r$,

- $\|\mathbf{x}_T\|_2 \leq R$;
- none of the matrices $\mathcal{P}_T \mathbb{E}[\mathbf{x} \mathbf{x}^T] \mathcal{P}_T$ are the zero matrix.

Following this notation, define

$$\phi_{\max} := \max_{|T| \leq k} \lambda_{\max}(\mathcal{P}_T \mathbf{C} \mathcal{P}_T) \quad \phi_{\min} := \min_{|T| \leq k} \lambda_{\max}(\mathcal{P}_T \mathbf{C} \mathcal{P}_T),$$

with $\lambda_{\max}(\cdot)$ the largest eigenvalue. The following corollary holds.

Proposition E.3 (Corollary 4 by (Elenberg et al., 2018)). *Consider a function l as in (7), and suppose that $\eta > 0$. Suppose that Assumption E.2 holds with parameters r, R , and suppose that the number of samples s is lower-bounded as*

$$s > \frac{R(\log r + r(1 + \log \frac{n}{k} - \log \delta))}{\phi_{\min}(1 + \varepsilon) \log(1 + \varepsilon) - \varepsilon}.$$

Then, with probability at least $1 - \delta$ the function l is (m, M) -(smooth, restricted concave), for all β with at most r non-zero coefficients. The parameters m and M are defined as $m = q(1 + \varepsilon)\phi_{\max} + \eta$ and $M = \eta$, with q a constant fulfilling $q \geq \max_i h^{-1}(\tau) Z''(\beta, \mathbf{x}_i)$ for $h^{-1}(\cdot), Z(\cdot, \cdot)$ as in (6).

Log-conditional with no assumptions on the regularization term. We now study log-conditional functions as in (6), with no additional assumption on the regularization term. For simplicity, we consider functions l as in (6) with $\eta = 0$. However, these results can easily be extended to the general case. The following lemma holds.

Lemma E.4 (Corollary 2 by (Elenberg et al., 2018)). Denote with r an upper-bound on the sparsity of the feature sets. Following the notation introduced above, suppose that the feature matrix \mathbf{X} consists of samples drawn from a sub-Gaussian distribution with parameter σ^2 and covariance matrix Σ . Then, for $\eta = 0$ the function l as in (7) is (M, m) -(smooth, restricted concave) with parameters

$$M = \alpha_u \lambda_{\max}(\Sigma) \left(3 + \frac{2nr}{s} \right) \quad \text{and} \quad m = \alpha_\ell - \frac{c^2 \sigma^2 k \log n}{\alpha_\ell s},$$

with high probability, for $s > 0$ sufficiently large. The constant α_ℓ depends on (σ^2, Σ) and k ; the constant α_u yields $\alpha_u \geq \max_i h^{-1}(\tau)^{-1} Z''(\beta, \mathbf{x}_i)$, with $h^{-1}(\cdot)$, $Z(\cdot, \cdot)$ as in (6).

E.2. Restricted Strong Concavity and Smoothness of the R^2 Objective

In this section, we prove guarantees for the R^2 objective. To this end, given a feature matrix consisting of n features and k observations, the R^2 objective with regularization can be written as

$$l(\beta) = R_{\mathbf{X}, \mathbf{y}}^2(\beta) = 1 - \frac{1}{k} \|\mathbf{y} - \langle \mathbf{X}, \beta \rangle\|_2^2. \quad (9)$$

The following lemma, similar to Lemma E.1, holds.

Proposition E.5. The R^2 objective l for the linear regression as in (9), is (m, M) -(restricted smooth, restricted strong concave) with parameters

$$m := \min_S \lambda_{\min}(\mathbf{X}_S^T \mathbf{X}_S) \quad \text{and} \quad M := \max_S \lambda_{\max}(\mathbf{X}_S^T \mathbf{X}_S).$$

This lemma can easily be extended to the case of an R^2 objective with regularization term.

F. The Computational Model

We assume access to an oracle that returns $\nabla l(\beta^{(\mathbb{T})})$, for a given input set \mathbb{T} . Elenberg et al. (2018) highlight the benefits of using this oracle model for feature selection, since access to the gradient is available from the inner optimization. In the case of a linear model, the gradient $\nabla l(\beta^{(\mathbb{T})})$ can be easily estimated for various metrics l expressing the goodness of fit. For instance, if l is the log-likelihood function, then the gradient can be computed in explicit form. For more complex models, stochastic lower-bounds of $\log \mathbb{P}_{\beta^{(s)}}(\mathbf{x})$ can be used, and then differentiated (Bamler et al., 2017; Nowozin, 2018). Similar considerations hold for other metrics, such as the R^2 objective (Elenberg et al., 2018). We also assume access to the independence oracle of the underlying p -system \mathcal{I} . The independence oracle takes as input a set \mathbb{T} , and it returns as output a Boolean value, true if $\mathbb{T} \in \mathcal{I}$ and false otherwise. This oracle is often assumed for optimizing functions over p -systems (Mirzasoleiman et al., 2016; Quinzan et al., 2021). Our method also works assuming access to a rank oracle, or a span oracle. We refer the reader to (Chekuri & Quanrud, 2019) for a description of these oracle models.

We evaluate performance using the notion of adaptivity. The adaptivity refers to the number of sequential rounds of the algorithm, wherein polynomial number of parallel queries are made in each round (Balkanski & Singer, 2018). Formally, given an oracle f , an algorithm is r -adaptive if every query q to the oracle f occurs at a round $i \in [r]$ such that q is independent of the answers $f(q')$ to all other queries q' at round i . This notion is closely related to the Parallel Random Access Machines (PRAM) model, as shown in Appendix G. We evaluate empirical speedup by the adaptivity of the oracle to evaluate $\nabla l(\beta^{(\mathbb{T})})$. We also evaluate the adaptivity of the independence oracle for the p -system \mathcal{I} .

G. Adaptivity and the PRAM Model

Recall that the adaptivity is defined as follows (Balkanski & Singer, 2018).

Definition G.1 (Adaptivity). Given an oracle f , an algorithm is r -adaptive if every query q to the oracle f occurs at a round $i \in [r]$ such that q is independent of the answers $f(q')$ to all other queries q' at round i .

The notion of adaptivity is closely related to other models such as Parallel Random Access Machines (PRAM). The PRAM model consists of a set of processors, that communicate via a single shared memory and a memory access unit. The adaptivity extend to PRAM via the notion of depth. The depth is the number of parallel steps in an algorithm or the longest chain of dependencies. We remark that the PRAM model assumes that the input is loaded in memory, whereas the adaptive complexity model only assumes access to an oracle function.

Algorithm 3 RNDSEQ (X, S)

```

A  $\leftarrow \emptyset$ ;
while  $X \neq \emptyset$  do
    Let  $\{x_1, \dots, x_n\}$  be a permutation of  $X$  chosen uniformly at random;
     $j^* \leftarrow \max\{j: S \cup A \cup \{x_1, \dots, x_i\}_{i \leq j} \in \mathcal{I}\}$ ;
     $A \leftarrow A \cup \{x_1, \dots, x_{j^*}\}$ ;
     $X \leftarrow \{e \in X \setminus (S \cup A): S \cup A \cup e \in \mathcal{I}\}$ ;
end while
return A;
    
```

H. The RNDSEQ Sub-Routine

The RNDSEQ sub-routine is presented in Algorithm 3. This algorithm correspond to Algorithm A by (Karp et al., 1988). This algorithm solves the problem of sampling a maximum independent set of \mathcal{I} uniformly at random. To our knowledge, no other algorithm is known for this problem, with better adaptivity than Algorithm 3. Given as input a ground set X , a current solution S , and a p -system \mathcal{I} , this algorithm finds a random set A such that $S \cup A$ is a maximum independent set for \mathcal{I} . This algorithm iteratively shuffles the set X , and then it identifies the longest prefix of this sequence that can be added to S , without violating side constraints. This prefix is then added to A . This algorithm terminates when $S \cup A \in \mathcal{I}$ is a maximal independent set.

This algorithm uses parallel calls to the independence oracle of \mathcal{I} . In fact, the evaluations for the feasibility of prefixes of A can be preformed in parallel. Hence, with this algorithm the adaptivity of the independence oracle corresponds to the number of iterations until convergence. It is well-known that this algorithm converges after expected $\mathcal{O}(\sqrt{r})$ iterations, with r the rank of \mathcal{I} (see Theorem 6 by (Karp et al., 1988)). This implies that the adaptivity of the independence oracle is $\mathcal{O}(\sqrt{r})$. Although it is not known if this upper-bound on the adaptivity is tight, it is known that there is no algorithm that finds a maximum independent set of \mathcal{I} with less then $\tilde{\Omega}(n^{1/3})$ rounds (see Theorem 7 by (Karp et al., 1988)).

I. Parameter Tuning for Algorithm 1

If l is the log-likelihood of a linear model, or the R^2 objective of a linear model with normalized response variable, then m and M can be estimated in terms of the design of the feature matrix (see Appendix E.1-E.2). These estimates need not be tight, and certifying bounds for m and M is NP-hard (Bandeira et al., 2013). In general, the constant m/M in Line 4 of Algorithm 1 requires tuning by making multiple runs of the algorithm. We remark that other known parallel algorithms for feature selection also require estimates of these parameters (Qian & Singer, 2019).

Parallel algorithms that estimate the rank are known for several p -systems. For instance, the rank of a graphic matroid can be estimated with parallel algorithms that compute spanning trees. Furthermore, parallel rank oracles are known for matroids that can be represented as independent sets of vectors in a given field (Borodin et al., 1982; Chistov, 1985; Ibarra et al., 1980; Mulmuley, 1987). These algorithms can also be used to estimate the rank of more complex constraints, such as the intersection of matroids or p -matchoids.

J. Proof of Theorem 4.1

We prove Theorem 4.1. The proof of this theorem is based on a few lemmas and propositions, which we discuss in Appendix J.1 before proving Theorem 4.1. On a high level, the proof of Theorem 4.1 is split into two separate cases. First we prove that Theorem 4.1 holds when Algorithm 1 terminates after ε^{-1} iterations of the outer While-loop of Algorithm 1. Then, we prove Theorem 4.1 under the assumption that Algorithm 1 finds a solution of size k . The first part of the proof is discussed in Appendix J.3 (see Theorem J.8), and the second case is discussed in Appendix J.2 (see Theorem J.5).

J.1. Preliminary Results

Our analysis is based on a few preliminary result, which we discuss in this section.

Theorem J.1. *Suppose the l is (M, m) -(smooth, strongly concave) on Ω_{2k} . Then, for each subsets $S, T \subseteq [p]$ of size at most*

k it holds

$$2M \sum_{j \in \mathbb{T}} f_S(j) \geq \|\nabla l(\boldsymbol{\beta}^{(S)})_{\mathbb{T}}\|_2^2 \geq 2m \sum_{j \in \mathbb{T}} f_S(j).$$

Proof. We start proving the first inequality. Fix a point $j \in \mathbb{T}$. Then, for any scalar α it holds

$$\begin{aligned} f_S(j) &= l(\boldsymbol{\beta}^{(S \cup \{j\})}) - l(\boldsymbol{\beta}^{(S)}) \geq l(\boldsymbol{\beta}^{(S)} + \alpha \mathbf{e}_j) - l(\boldsymbol{\beta}^{(S)}) && \text{(maximality of } \boldsymbol{\beta}^{(S \cup \{j\})}) \\ &\geq \langle \nabla l(\boldsymbol{\beta}^{(S)}), \alpha \mathbf{e}_j \rangle - \frac{M}{2} \alpha^2, && \text{(restricted smoothness)} \end{aligned} \quad (10)$$

By substituting $\alpha = \frac{1}{M} \langle \nabla l(\boldsymbol{\beta}^{(S)}), \mathbf{e}_j \rangle$ in (10), we get

$$2M \sum_{j \in \mathbb{T}} f_S(j) \geq \sum_{j \in \mathbb{T}} \left\langle \nabla l(\boldsymbol{\beta}^{(S)}), \mathbf{e}_j \right\rangle^2 = \|\nabla l(\boldsymbol{\beta}^{(S)})_{\mathbb{T}}\|_2^2,$$

and the first inequality follows. To conclude the proof, note that it holds

$$\begin{aligned} f_S(j) &= l(\boldsymbol{\beta}^{(S \cup \{j\})}) - l(\boldsymbol{\beta}^{(S)}) \\ &\leq \langle \nabla l(\boldsymbol{\beta}^{(S)}), \boldsymbol{\beta}^{(S \cup \{j\})} - \boldsymbol{\beta}^{(S)} \rangle - \frac{m}{2} \|\boldsymbol{\beta}^{(S \cup \{j\})} - \boldsymbol{\beta}^{(S)}\|_2^2 && \text{(restricted strong concavity)} \\ &\leq \max_{\mathbf{v}: \mathbf{v}_{(S \cup \{j\})=0}} \langle \nabla l(\boldsymbol{\beta}^{(S)}), \mathbf{v} - \boldsymbol{\beta}^{(S)} \rangle - \frac{m}{2} \|\mathbf{v} - \boldsymbol{\beta}^{(S)}\|_2^2. && \text{(maximality of } \mathbf{v}) \end{aligned} \quad (11)$$

By setting $\mathbf{v} = \boldsymbol{\beta}^{(S)} + \frac{1}{m} \langle \nabla l(\boldsymbol{\beta}^{(S)}), \mathbf{e}_j \rangle$ in (11) we get

$$\frac{\|\nabla l(\boldsymbol{\beta}^{(S)})_{S^*}\|_2^2}{2m} \geq l(\boldsymbol{\beta}^{(S \cup \{j\})}) - l(\boldsymbol{\beta}^{(S)}) = f_S(j).$$

By taking the sum over all $j \in \mathbb{T}$ and rearranging we get

$$\|\nabla l(\boldsymbol{\beta}^{(S)})_{\mathbb{T}}\|_2^2 = \sum_{j \in \mathbb{T}} \frac{\|\nabla l(\boldsymbol{\beta}^{(S)})_{S^*}\|_2^2}{2m} \geq \sum_{j \in \mathbb{T}} f_S(j),$$

and the claim follows. \square

In our analysis, we also use the following well-known result.

Theorem J.2 (Theorem 1 by (Elenberg et al., 2018)). *Suppose the l is (M, m) -(smooth, strongly concave) on Ω_{2k} . Then, for each subsets $S, \mathbb{T} \subseteq [p]$ of size at most k it holds*

$$2M f_S(\mathbb{T}) \geq \|\nabla l(\boldsymbol{\beta}^{(S)})_{\mathbb{T}}\|_2^2 \geq 2m f_S(\mathbb{T}).$$

By combining Theorem J.1 with Theorem J.2, we get the following corollary.

Corollary J.3. *Suppose the l is (M, m) -(smooth, strongly concave) on Ω_{2k} . Then, for each subsets $S, \mathbb{T} \subseteq [p]$ of size at most k it holds*

$$\frac{M}{m} f_S(\mathbb{T}) \geq \sum_{j \in \mathbb{T}} f_S(j) \geq \frac{m}{M} f_S(\mathbb{T}).$$

We also make use of the following technical proposition.

Proposition J.4 (Proposition 2.2. by (Nemhauser et al., 1978)). *Let $\{x_1, \dots, x_m\}$ and $\{y_1, \dots, y_m\}$ be two sequences of non-negative real numbers. Suppose that it holds $\sum_j x_j \leq 1$ for all $j \in [m]$. Then,*

$$\sum_{j=1}^m y_j \geq \sum_{j=1}^m x_j y_j.$$

J.2. If Algorithm 1 Outputs a Maximum Independent Set

We now prove Theorem 4.1, assuming that Algorithm 1 outputs a solution S of maximum size, before performing ε^{-1} iterations of the outer While-loop of Algorithm 1. Formally, we prove the following theorem.

Theorem J.5. *Define the function f as in (1), with a log-likelihood function that is (M, m) -(smooth, strongly concave) on Ω_{2r} . Suppose that Algorithm 1 outputs a solution S^* such that $\text{Cond}(S^*) = \emptyset$. Then,*

$$\frac{\mathbb{E}[f(S^*)]}{\text{OPT}} \geq \frac{1}{1+p} \left(1 - \exp \left\{ -(1-\varepsilon)^2 \frac{m^3}{M^3} \right\} \right),$$

for all $0 < \varepsilon < 1$. Furthermore, in the specific case when \mathcal{I} is r -sparsity constraint over $[n]$, then,

$$\frac{\mathbb{E}[f(S^*)]}{\text{OPT}} \geq 1 - \exp \left\{ -(1-\varepsilon)^2 \frac{m^2}{M^2} \right\}.$$

The proof of Theorem J.5 is based on the following two additional lemma.

Lemma J.6. *At any point during the optimization process it holds*

$$(1-\varepsilon)^{-1} \frac{rM}{m} t \geq 2m (\text{OPT} - f(S)),$$

with S the current solution.

Proof. Denote with \bar{S} a solution of size at most $|\bar{S}| \leq r$ maximizing $f(S \cup \bar{S})$, and let $T \subseteq [n]$ be a set maximizing $\|\nabla l(\beta^{(S)})_T\|_2^2$, such that $|T| \leq r$ and $S \cup \{s\} \in \mathcal{I}$ for all $s \in T$. Note that it holds

$$(1-\varepsilon)^{-1} \frac{rM}{m} t = \frac{r}{|T|} \|\nabla l(\beta^{(S)})_T\|_2^2 \geq \|\nabla l(\beta^{(S)})_T\|_2^2, \quad (12)$$

where the first inequality follows by the definition of t , and the second one follows since $|T| \leq r$. We first prove the claim when t is updated at the beginning of each iteration of the outer While-loop of Algorithm 1. It holds

$$\begin{aligned} & l(\beta^{(S \cup \bar{S})}) - l(\beta^{(S)}) \\ & \leq \langle \nabla l(\beta^{(S)}), \beta^{(S \cup \bar{S})} - \beta^{(S)} \rangle - \frac{m}{2} \|\beta^{(S \cup \bar{S})} - \beta^{(S)}\|_2^2 \quad (\text{restricted strong concavity}) \\ & \leq \max_{\mathbf{v}: \mathbf{v}_{(S \cup \bar{S}) \neq 0}} \langle \nabla l(\beta^{(S)}), \mathbf{v} - \beta^{(S)} \rangle - \frac{m}{2} \|\mathbf{v} - \beta^{(S)}\|_2^2. \quad (\text{maximality of } \mathbf{v}) \end{aligned} \quad (13)$$

By setting $\mathbf{v} = \frac{1}{m} \beta^{(S)} + \nabla l(\beta^{(S)})_{\bar{S}}$ in the inequality above we get

$$\begin{aligned} f(\text{OPT}) - f(S) & = l(\beta^{(S \cup \bar{S})}) - l(\beta^{(S)}) \quad (\text{maximality of } \bar{S} \text{ and monotonicity}) \\ & \leq \frac{\|\nabla l(\beta^{(S)})_{\bar{S}}\|_2^2}{2m} \quad (\text{substituting } \mathbf{v} \text{ in (13)}) \\ & \leq \frac{\|\nabla l(\beta^{(S)})_T\|_2^2}{2m}. \quad (\text{maximality of } T) \end{aligned} \quad (14)$$

The claim follows by combining (14) and (12). Suppose now that the current solution S is updated to S' during the inner While-loop of Algorithm 1. Then, $f(S) \geq f(S')$ due to monotonicity, and the claim holds. \square

Lemma J.7. *At any given time step, suppose that the current solution S is updated to $S \cup \{a_1, \dots, a_{j^*}\}$, and define $S_j = S \cup \{a_1, \dots, a_j\}$ for all $j \in [j^*]$. Then it holds*

$$\mathbb{E}[f_{S_{j-1}}(S_j)] \geq (1-\varepsilon)^2 \frac{m^2}{rM^2} (\text{OPT} - \mathbb{E}[f(S_{j-1})]).$$

Proof. Fix all random decisions of Algorithm 1 until the point a_j is added to the current solution. Define the sets $\mathcal{X}_j^{\mathcal{I}} := \{e \in \mathcal{X} \setminus \{a_1, \dots, a_{j-1}\} : \{a_1, \dots, a_{j-1}\} \cup \{a\} \in \mathcal{I}\}$. Then, it holds

$$\begin{aligned} \mathbb{E}_{a_j} [f_{S_{j-1}}(a_j)] &\geq \frac{1}{2M} \mathbb{E}_{a_j} [\langle \nabla l(\beta^{(S_{j-1})}), \mathbf{e}_{a_j} \rangle^2] && \text{(Theorem J.1)} \\ &\geq \frac{1}{2M} \mathbb{P}_{a_j}(\langle \nabla l(\beta^{(S_{j-1})}), \mathbf{e}_{a_j} \rangle^2 \geq t)t && \text{(Markov's inequality)} \end{aligned} \quad (15)$$

By design of the RNDSEQ subroutine, each point a_j is sampled uniformly at random from the set $\mathcal{X}_j^{\mathcal{I}}$, i.e. $a_j \sim \mathcal{U}(\mathcal{X}_j^{\mathcal{I}})$. Hence, it holds $\mathbb{P}_{a_j}(\langle \nabla l(\beta^{(S_{j-1})}), \mathbf{e}_{a_j} \rangle^2 \geq t) = |\mathcal{X}_{j-1}| / |\mathcal{X}_j^{\mathcal{I}}|$. Combining this observation with (15) we get

$$\begin{aligned} \mathbb{E}_{a_j} [f_{S_{j-1}}(a_j)] &\geq \frac{1}{2M} \frac{|\mathcal{X}_{j-1}|}{|\mathcal{X}_j^{\mathcal{I}}|} t, && \text{(by the sub-sampling procedure)} \\ &\geq \frac{1}{2M} \frac{|\mathcal{X}_{j-1}|}{|\mathcal{X}|} t, && \text{(since } \mathcal{X}_j^{\mathcal{I}} \subseteq \mathcal{X}\text{)} \\ &\geq (1 - \varepsilon) \frac{1}{2M} t && \text{(since } |\mathcal{X}_{j-1}| \geq (1 - \varepsilon) |\mathcal{X}|\text{)} \\ &\geq (1 - \varepsilon)^2 \frac{m^2}{rM^2} (\text{OPT} - f(S_{j-1})), && \text{(by Lemma J.6)} \end{aligned}$$

The claim follows by taking the expectation on both sides. \square

Using this lemma, we can now prove Theorem J.5.

Proof of Theorem J.5. Denote with $\{a_1, \dots, a_j\}$ the first j points added to the solution S^* , sorted in the order that they were added to it, and define the constant $c := (1 - \varepsilon)^2 m^2 / rM^2$. Using an induction argument on j , we prove that it holds

$$\frac{\mathbb{E} [f(\{a_1, \dots, a_j\})]}{\text{OPT}} \geq (1 - (1 - c)^j). \quad (16)$$

The base case with $j = 0$ holds, due to the non-negativity of the function f . For the inductive case, we have that it holds

$$\begin{aligned} \mathbb{E} [f(\{a_1, \dots, a_j\})] &\geq \mathbb{E} [f(\{a_1, \dots, a_{j-1}\})] + c(\text{OPT} - f(\{a_1, \dots, a_{j-1}\})) && \text{(Lemma J.7)} \\ &\geq (1 - c) (1 - (1 - c)^{j-1}) \text{OPT} + c\text{OPT} && \text{(by induction)} \\ &\geq (1 - (1 - c)^j) \text{OPT}, \end{aligned}$$

and (16) holds. It follows that for $j = |S^*|$ we have

$$\frac{\mathbb{E} [f(S^*)]}{\text{OPT}} \geq 1 - (1 - c)^{|S^*|} \geq 1 - \exp \left\{ -\frac{|S^*|}{k} c \right\}.$$

In the case of a r -sparsity constraint we have that $|S^*| = r$, and the claim follows. In the case of a p -system constraint, we have that $|S^*| \geq pr$, hence

$$\frac{\mathbb{E} [f(S^*)]}{\text{OPT}} \geq 1 - \exp \left\{ -(1 - \varepsilon)^2 p \frac{m^2}{M^2} \right\} \geq \frac{1}{1 + p} \left(1 - \exp \left\{ -(1 - \varepsilon)^2 \frac{m^3}{M^3} \right\} \right),$$

and the claim also holds. \square

J.3. If Algorithm 1 Terminates after ε^{-1} Iterations

We now prove Theorem 4.1, assuming that the FAST_{OMP} terminates after ε^{-1} iterations of the outer While-loop of Algorithm 1. Specifically, we prove the following theorem.

Theorem J.8. Define the function f as in (1), with a log-likelihood function that is (M, m) -(smooth, strongly concave) on Ω_{2r} . Suppose that Algorithm 1 terminates after ε^{-1} iterations of the outer-While loop. Then,

$$\frac{\mathbb{E} [f(\mathbf{S}^*)]}{\text{OPT}} \geq \frac{1}{1+p} \left(1 - \exp \left\{ -(1-\varepsilon)^2 \frac{m^3}{M^3} \right\} \right),$$

for all $0 < \varepsilon < 1$. Furthermore, in the specific case when \mathcal{I} is r -sparsity constraint over $[n]$, then,

$$\frac{\mathbb{E} [f(\mathbf{S}^*)]}{\text{OPT}} \geq 1 - \exp \left\{ -(1-\varepsilon)^2 \frac{m^2}{M^2} \right\}.$$

In order to prove this theorem, we introduce additional notation. We denote with \mathbf{S}_i the current solution at the beginning of the i -th iteration of the outer While-loop of Algorithm 1. Furthermore, denote with $\bar{\mathbf{S}} \subseteq \mathbf{T} \setminus \mathbf{S}_i$ a feasible set, such that $f(\bar{\mathbf{S}} \cup \mathbf{S}_i) = \text{OPT}$, and denote with a_j the j -th element added to the solution \mathbf{S} . The proof of this theorem is based on the following lemma.

Lemma J.9. It holds

$$\frac{M}{m} f_{\mathbf{S}_i}(\mathbf{S}_{i+1}) + \sum_{e \in \bar{\mathbf{S}} \setminus \text{Cond}(\mathbf{S}_i)} \varepsilon f_{\mathbf{S}_i}(e) \geq \varepsilon \frac{m}{M} f_{\mathbf{S}_i}(\bar{\mathbf{S}}).$$

Proof. Fix all random decision of Algorithm 1, up to the $(i+1)$ -th iteration of the outer While-loop of Algorithm 1. Let $\mathbf{T} \subseteq [n]$ be a set maximizing $\|\nabla l(\boldsymbol{\beta}^{(\mathbf{S}_i)})_{\mathbf{T}}\|_2^2$, such that $|\mathbf{T}| \leq r$ and $\mathbf{T} \subseteq \mathcal{I}$. Due to the assumption on the stopping criterion, it holds $\mathbf{X} = \emptyset$ at the end of iteration i . This means that each point $j \in (\mathbf{T} \setminus \mathbf{S}_i) \cap \text{Cond}(\mathbf{S}_i)$ was discarded at some point during the previous iteration. Denote with \mathbf{U}_j the current solution when j was discarded. Then, it holds

$$(1-\varepsilon) \frac{2m}{r} \sum_{e \in \mathbf{T}} f_{\mathbf{S}_i}(e) \geq (1-\varepsilon) \frac{m}{rM} \|\nabla l(\boldsymbol{\beta}^{(\mathbf{S}_i)})_{\mathbf{T}}\|_2^2 = t \quad (17)$$

where the first inequality follows by Theorem J.1, and the second one follows by the definition of t . Since the point j was discarded and since $j \in \text{Cond}(\mathbf{S}_i)$, then it must hold $t \geq \langle \nabla l(\boldsymbol{\beta}_j^{(\mathbf{U}_j)}), \mathbf{e}_j \rangle^2$. Note also that by the RSC/RSM properties of the function l it holds $\langle \nabla l(\boldsymbol{\beta}^{(\mathbf{U}_j)}), \mathbf{e}_j \rangle^2 \geq \langle \nabla l(\boldsymbol{\beta}^{(\mathbf{S}_{i+1})}), \mathbf{e}_j \rangle^2$ (Elenberg et al., 2018). Combining these observations with (17) we get

$$(1-\varepsilon) \frac{2m}{r} \sum_{e \in \mathbf{T}} f_{\mathbf{S}_i}(e) \geq \langle \nabla l(\boldsymbol{\beta}^{(\mathbf{S}_{i+1})}), \mathbf{e}_j \rangle^2 \geq 2m f_{\mathbf{S}_{i+1}}(j). \quad (18)$$

By taking the sum over all points $j \in (\mathbf{T} \setminus \mathbf{S}_i) \cap \text{Cond}(\mathbf{S}_i)$ and rearranging, we get

$$\begin{aligned} (1-\varepsilon) \sum_{e \in \mathbf{T}} f_{\mathbf{S}_i}(e) &\geq (1-\varepsilon) \frac{|(\mathbf{T} \setminus \mathbf{S}_i) \cap \text{Cond}(\mathbf{S}_i)|}{r} \sum_{e \in \mathbf{T}} f_{\mathbf{S}_i}(e) && \text{(by definition of } r) \\ &\geq \sum_{j \in (\mathbf{T} \setminus \mathbf{S}_i) \cap \text{Cond}(\mathbf{S}_i)} \bar{f}_{\mathbf{S}_{i+1}}(j) && \text{(it follows from (18))} \end{aligned} \quad (19)$$

By rearranging (19) we get

$$\begin{aligned} \sum_{e \in \mathbf{S}_{i+1}} f_{\mathbf{S}_i}(e) &\geq \sum_{e \in (\mathbf{T} \setminus \mathbf{S}_i) \cap \text{Cond}(\mathbf{S}_i)} \varepsilon f_{\mathbf{S}_i}(e) \\ &\geq \sum_{e \in \bar{\mathbf{S}} \cap \text{Cond}(\mathbf{S}_i)} \varepsilon f_{\mathbf{S}_i}(e) && \text{(by the definition of } \bar{\mathbf{S}}) \\ &\geq \sum_{e \in \bar{\mathbf{S}}} \varepsilon f_{\mathbf{S}_i}(e) - \sum_{e \in \bar{\mathbf{S}} \setminus \text{Cond}(\mathbf{S}_i)} \varepsilon f_{\mathbf{S}_i}(e) && \text{(by linearity)} \end{aligned} \quad (20)$$

Hence, it holds

$$\begin{aligned}
 \frac{M}{m} f_{S_i}(S_{i+1}) &\geq \sum_{e \in S_{i+1}} f_{S_i}(e) && \text{(Corollary J.3)} \\
 &\geq \sum_{e \in \bar{S}} \varepsilon f_{S_i}(e) - \sum_{e \in \bar{S} \setminus \text{Cond}(S_i)} \varepsilon f_{S_i}(e) && \text{(it follows from (20))} \\
 &\geq \varepsilon \frac{m}{M} f_{S_i}(\bar{S}) - \sum_{e \in \bar{S} \setminus \text{Cond}(S_i)} \varepsilon f_{S_i}(e). && \text{(Corollary J.3)}
 \end{aligned}$$

The claim follows by rearranging. \square

In order to continue with the proof, we also use the following lemma.

Lemma J.10. *It holds*

$$pf(S_i) \geq (1 - \varepsilon) \frac{m^2}{M^2} \sum_{e \in \bar{S} \setminus \text{Cond}(S_i)} f_{S_i}(e).$$

Proof. Fix an index $j \leq |S_i|$, and fix all random decisions of Algorithm 1 until the point a_j is added to the current solution. For each element a_j , define the set $D_j := (\bar{S} \cap \text{Cond}(\{a_1, \dots, a_{j-1}\})) \setminus (\bar{S} \cap \text{Cond}(\{a_1, \dots, a_j\}))$. Note that these sets consist of all points in \bar{S} that yield a feasible solution when added to $\{a_1, \dots, a_{j-1}\}$, but that violate side constraints when added to $\{a_1, \dots, a_j\}$. Note also that it holds $D_1 \cup \dots \cup D_j = \bar{S} \setminus \text{Cond}(\{a_1, \dots, a_j\})$. Furthermore, define the sets

$$\mathcal{X}_j^{\mathcal{I}} := \{e \in \mathcal{X} \setminus \{a_1, \dots, a_{j-1}\} : \{a_1, \dots, a_{j-1}\} \cup \{a\} \in \mathcal{I}\}.$$

Then, it holds

$$\begin{aligned}
 \mathbb{E}_{a_j} [f_{\{a_1, \dots, a_{j-1}\}}(a_j)] &\geq \frac{1}{2M} \mathbb{E}_{a_j} [\langle \nabla l(\beta^{\{a_1, \dots, a_{j-1}\}}), \mathbf{e}_{a_j} \rangle^2] && \text{(by Lemma J.1)} \\
 &\geq \frac{1}{2M} \mathbb{P}_{a_j}(\langle \nabla l(\beta^{\{a_1, \dots, a_{j-1}\}}), \mathbf{e}_{a_j} \rangle^2 \geq t) t && \text{(by Markov)} \quad (21)
 \end{aligned}$$

Note that the RNDSEQ subroutine samples points a_j uniformly at random $a_j \sim \mathcal{U}(\mathcal{X}_j^{\mathcal{I}})$. Hence,

$$\mathbb{P}_{a_j}(\langle \nabla l(\beta^{(S_{j-1})}), \mathbf{e}_{a_j} \rangle^2 \geq t) = \frac{|\mathcal{X}_{j-1}|}{|\mathcal{X}_j^{\mathcal{I}}|} \geq \frac{|\mathcal{X}_{j-1}|}{|\mathcal{X}|}, \quad (22)$$

where the last inequality holds, since $\mathcal{X}_j^{\mathcal{I}} \subseteq \mathcal{X}$. Combining this observation with (21) we get

$$\begin{aligned}
 \mathbb{E}_{a_j} [f_{\{a_1, \dots, a_{j-1}\}}(a_j)] &\geq \frac{1}{2M} \frac{|\mathcal{X}_j|}{|\mathcal{X}|} t && \text{(it follows by (22))} \\
 &\geq \frac{1 - \varepsilon}{2M} t && (|\mathcal{X}_{j-1}| \geq (1 - \varepsilon) |\mathcal{X}|) \\
 &= \frac{1 - \varepsilon}{2M^2} \frac{m}{|\mathcal{T}|} \|\nabla l(\beta^{\{a_1, \dots, a_{j-1}\}})_{\mathcal{T}}\|_2^2 && \text{(by the definition of } t) \\
 &\geq \frac{1 - \varepsilon}{2M^2} \frac{m}{|D_j|} \|\nabla l(\beta^{\{a_1, \dots, a_{j-1}\}})_{D_j}\|_2^2, && (\mathcal{T} \text{ is maximal})
 \end{aligned}$$

By taking the expected value on both sides in the chain of inequalities above, we get

$$\sum_{j \leq |S_i|} |D_j| \mathbb{E}_{a_j} [f_{\{a_1, \dots, a_{j-1}\}}(a_j)] \geq (1 - \varepsilon) \frac{m}{2M^2} \sum_{j \leq |S_i|} \mathbb{E} [\|\nabla l(\beta^{(S_i)})_{D_j}\|_2^2]. \quad (23)$$

In order to continue with the proof, we give an upper-bound on the size of the sum $\sum_j |D_j|$. To this end, note that the set S_i is a maximum independent set over the ground set

$$S_i \cup (D_1 \cup \dots \cup D_{|S_i|}) = S_i \cup (\bar{S} \setminus \text{Cond}(S_i)).$$

In fact, the set S_i is independent by definition, and that any point $s \in (\bar{S} \setminus \text{Cond}(S_i)) \setminus S_i$ yields $S_i \cup \{s\} \notin \mathcal{I}$. Hence S_i is a maximum independent set as claimed. Note also that $D_1 \cup \dots \cup D_{|S_i|} \subseteq \bar{S}$ is an independent set, due to the subset-closure of \mathcal{I} . Since \mathcal{I} is a p -system, then it holds

$$|D_1| + \dots + |D_{|S_i|}| = |D_1 \cup \dots \cup D_{|S_i|}| \leq p |S_i|. \quad (24)$$

Hence, it holds

$$\begin{aligned} p \sum_{j \leq |S_i|} \mathbb{E}_{a_j} [f_{\{a_1, \dots, a_{j-1}\}}(a_j)] &\geq \sum_{j \leq |S_i|} |D_j| \mathbb{E}_{a_j} [f_{\{a_1, \dots, a_{j-1}\}}(a_j)] && \text{(it follows by (24))} \\ &\geq (1 - \varepsilon) \frac{m}{2M^2} \sum_{j \leq |S_i|} \mathbb{E} \left[\|\nabla l(\beta^{(S_i)})_{D_j}\|_2^2 \right] && \text{(it follows by (23))} \\ &\geq (1 - \varepsilon) \frac{m^2}{M^2} \sum_{e \in D_1 \cup \dots \cup D_{|S_i|}} f_{S_i}(e) && \text{(by Theorem J.1).} \end{aligned}$$

The claim follows since $D_1 \cup \dots \cup D_{|S_i|} = \bar{S} \setminus \text{Cond}(S_i)$. \square

We now have all necessary tools to prove Theorem J.8.

Proof of Theorem J.8. We first prove the claim, assuming that \mathcal{I} is a general p -system. In this case, by combining Lemma J.9 with Lemma J.10 it holds

$$\frac{M}{m} \mathbb{E} [f_{S_i}(S_{i+1})] + \varepsilon p \frac{M^2}{m^2} \mathbb{E} [f(S_i)] \geq \varepsilon (1 - \varepsilon) \frac{m}{M} \mathbb{E} [f_{S_i}(\bar{S})]. \quad (25)$$

To continue, define the constant $c = (1 - \varepsilon)m^3/M^3$. We prove by induction on i that it holds

$$(1 + \varepsilon i p) \mathbb{E} [f(S_i)] \geq (1 - (1 - \varepsilon c)^i) \text{OPT}. \quad (26)$$

The base case with $S_0 = \emptyset$ trivially follows, since the function f is non-negative. For the inductive case, suppose that the claim holds for $\mathbb{E} [f(S_{i-1})]$. Then,

$$\begin{aligned} (1 + \varepsilon i p) \mathbb{E} [f(S_i)] &\geq \mathbb{E} [f(S_i)] + \varepsilon i p \mathbb{E} [f(S_{i-1})] && \text{(by monotonicity)} \\ &\geq \mathbb{E} [f(S_{i-1})] + \varepsilon c \mathbb{E} [f_{S_i}(\bar{S})] + \varepsilon (i - 1) p \mathbb{E} [f(S_{i-1})] && \text{(it follows by (25))} \\ &\geq (1 - \varepsilon c) \mathbb{E} [f(S_{i-1})] + \varepsilon c \text{OPT} + \varepsilon (i - 1) p \mathbb{E} [f(S_{i-1})] && \text{(by monotonicity)} \\ &\geq (1 - \varepsilon c) (1 - (1 - \varepsilon c)^{i-1}) \text{OPT} + \varepsilon c \text{OPT} && \text{(by induction)} \\ &\geq (1 - (1 - \varepsilon c)^i) \text{OPT}. \end{aligned}$$

Hence, (26) holds. It follows that

$$\begin{aligned} \mathbb{E} [f(S^*)] &= \mathbb{E} [f(S_{\lfloor 1/\varepsilon \rfloor})] && \text{(by the stopping criterion)} \\ &\geq \frac{1}{1 + \varepsilon \lfloor 1/\varepsilon \rfloor p} \left(1 - \left(1 - \varepsilon (1 - \varepsilon) \frac{m^3}{M^3} \right)^{\lfloor 1/\varepsilon \rfloor} \right) \text{OPT} && \text{(it follows by (26))} \\ &\geq \frac{1}{1 + p} \left(1 - \exp \left\{ - (1 - \varepsilon) \frac{m^3}{M^3} \right\} \right) \text{OPT}, \end{aligned}$$

as claimed.

We conclude by proving the claim in the special case that \mathcal{I} is a r -sparsity constraint. Since the algorithm terminates before a solution of size r is found, we have that $\text{Cond}(S_i) = [n] \setminus S_i$ for all iterations i . Hence, $D_1 \cup \dots \cup D_i = \emptyset$ and Lemma J.9 yields

$$\mathbb{E} [f_{S_i}(S_{i+1})] \geq \varepsilon \frac{m^2}{M^2} \mathbb{E} [f_{S_i}(S^*)].$$

With this inequality, we can use an inductive argument similar to the proof for the general case, and obtain an improved lower-bound on the solution quality. \square

K. Proof of Theorem 4.2

We conclude by giving upper-bounds on the run time and adaptivity for Algorithm 1. Recall that the notion of adaptivity is given in Definition G.1.

In order to prove Theorem 4.2, we use the following well-known estimate on the number of adaptive rounds of the RNDSEQ sub-routine (see Appendix H).

Theorem K.1 (Theorem 6 by (Karp et al., 1988)). *Algorithm H terminates after expected $\mathcal{O}(\sqrt{r})$ steps, with r the rank of the independent system \mathcal{I} .*

Note that this theorem implies that the number of adaptive rounds of the independence oracle for Algorithm H is $\mathcal{O}(\sqrt{r})$ in expected value. In fact, in each step of Algorithm H, queries to the independence oracle can be performed in parallel. Using this result, we can now prove Theorem 4.2.

Proof of Theorem 4.2. We first give upper-bounds for the oracle function that accesses $\nabla l(\cdot)$. To this end, observe that there are two While-loops in Algorithm 1. The outer while-loop terminates after at most ε^{-1} iteration. The inner While-loop terminates after $\mathcal{O}(\varepsilon^{-1} \log n)$ iterations, since at each iteration the size of X decreases at least of a multiplicative factor of $1 - \varepsilon$. Hence, the rounds of calls to the oracle function is $\mathcal{O}(\varepsilon^{-2} \log n)$. Furthermore, at each iteration of the inner While-loop, at most r parallel calls to $\nabla l(\cdot)$ are performed. It follows that the total number of oracle calls is $\mathcal{O}(\varepsilon^{-2} r \log n)$.

We now estimate the number of adaptive rounds and run time for the calls to the independence oracle. To this end, note that is oracle is called by the RNDSEQ sub-routine, and it is also evaluated nr times in parallel during the inner While-loop of Algorithm 1. From Theorem K.1 it follows that the number of adaptive rounds is $\mathcal{O}(\varepsilon^{-2} \sqrt{r} \log n)$, and that the total number of calls to the oracle function is $\mathcal{O}(\varepsilon^{-2} nr \log n)$ as claimed. \square

L. Benchmarks²

In this section, we describe the benchmark algorithms we use for comparison against the proposed framework. Our goal in empirical evaluation is to illustrate the accuracy vs speedup tradeoff that follows from our computational model of adaptive sampling based matching pursuit under constraints. The main purpose of using a technique like adaptive sampling is to obtain speedups by reducing the number of oracle evaluations compared to their non-adaptive counterparts, with some admissible loss in accuracy (Balkanski et al., 2019). Indeed, there are many other feature selection algorithms. We choose feature selection benchmarks that help us illustrate the said speedup obtained when using our method while ensuring graceful degradation in accuracy in simulated and real-world use-cases. These algorithms include popular selection methods (Kalimeris et al., 2019; Krause & Cevher, 2010; Elenberg et al., 2018; Grgic-Hlaca et al., 2018b). Other popular algorithms for feature selection include the Maximum Relevance Minimum Redundancy (mRMR) (Zhao et al., 2019; Ding & Peng, 2005) and the Conditional Mutual Information Maximisation (CMIM) filters (Torkkola, 2003) among others. These algorithms iteratively adds features by maximizing suitable objectives, such as scores based on the F-statistic, or Mutual Information gain. However, these algorithms for general p -system side constraints require $\Omega(r)$ sequential calls to the independence oracle for a solution size of r . Thus, such methods are impractical and will be trivially too slow compared to our method. We consider the following algorithms to compare against:

SDS_{MA}: Starting from the empty set, this algorithm adds feasible points to the current solution in a greedy fashion (Krause & Cevher, 2010; Elenberg et al., 2018). This algorithm uses oracle access to the function f as in (1).

SDS_{OMP}: Starting from the empty set, this algorithm iteratively adds a feature s to the current solution S if it maximizes the dot product $\langle \nabla l(\beta^{(S)}), e_s \rangle$ (Krause & Cevher, 2010; Elenberg et al., 2018).

DASH: This algorithm follows a computationally similar model, and achieves strong approximation guarantees on the subset selection problem, under the RSC/RSM assumption (Qian & Singer, 2019). It also uses oracle access to the function f but can only handle r -sparsity constraints.

ISK: This algorithm is the iterated submodular-cost knapsack algorithm proposed by Iyer & Bilmes (2013). It was used by Grgic-Hlaca et al. (2018b) to perform feature selection on the ProPublica COMPAS dataset. The ISK can only handle r -sparsity constraints, and it has no known guarantees for the problem (2).

²We parallelize all algorithms for fair comparison.

Lasso: We also compare against the popular Lasso regression. Tuning Lasso to obtain a solution of a desired size is hard (Mairal & Yu, 2012). In our experiments, we vary the parameter manually and benchmark against the resulting solution size.

Random: This simple algorithm outputs a maximum independent set of \mathcal{I} chosen uniformly at random. We use the RNDSEQ sub-routine to generate this set (see Appendix H).

M. Feature selection on Non-volatile Memory (NVM)

The emergence of CPU-attached persistent memory technology, such as Intel’s Optane Non-Volatile Memory (NVM), has opened opportunities for running large datasets on a single server. A single server may have up to 6TB of NVM, that can be accessed through the memory bus. The use of NVM compared to DRAM is beneficial for two reasons. First, due to its large capacity, it is suitable to handle large datasets. Secondly, as NVM is non-volatile, it can easily support fault-tolerance and recovery of the computation when a server crashes. Furthermore, even though NVM latency is slightly slower than DRAM, it has much better latency than SSD of orders of magnitude.

For running our Python code on NVM, we used MCAS (Waddington et al., 2021b;a) (Memory Centric Active Storage), which is an advanced client-server in-memory object store designed from the ground up to leverage persistent memory. MCAS supports “pushdown” operations on the client-side which are termed Active Data Objects (ADO), and has a Python plugin that allows zero-copy for Numpy data access. Our benchmark for evaluating feature selection with persistent memory is done by integrating the algorithms to MCAS as an ADO using the above mentioned Python plugin. We use an Intel Xeon Gold 6248 server with 80 CPUs at 2.50GHz. The server is equipped with 384GB DDR4 DRAM and 1512GB Optane DC.

By experimenting with the same datasets as in Section 5 with the SDS_{OMP} and FAST_{OMP} , we observed no significant variation in run time compared to our DRAM implementation. This is due to the fact that the time required to copy data from NVM to DRAM is negligible, in comparison with the compute time of training phases. However, when performing tasks on the DRAM that use more memory than the DRAM capacity, we might observe a significant decrease in the performance. For this reason, in future work we intend to investigate the trade-off between training time on different media (NVM versus DRAM and SSD) of tasks that use more memory than the DRAM capacity.

N. The ProPublica COMPAS Dataset

The ProPublica COMPAS dataset was constructed in 2016, using data of defendants from Broward County, FL, who had been arrested in 2013 or 2014 and assessed with the COMPAS risk screening system. ProPublica then collected data on future arrests for these defendants through the end of March 2016, in order to study how the COMPAS score predicted recidivism (Angwin et al., 2016). Based on its analysis, ProPublica concluded that the COMPAS risk score was racially biased (Berk et al., 2021).

The ProPublica COMPAS data has become one of the key bench-marking datasets for testing algorithmic fairness definitions and procedures (Corbett-Davies et al., 2017; Chouldechova, 2017; Corbett-Davies et al., 2017; Cowgill & Tucker, 2019; Rudin et al., 2018; Zafar et al., 2017b;c). However, (Bao et al., 2021a) notes that there are inaccuracies in the COMPAS dataset. For instance, COMPAS race categories lack Native Hawaiian or Other Pacific Islander, and it redefines Hispanic as race instead of ethnicity.

This dataset consists of the following features: “number of prior criminal offenses”, “arrest charge description”, “charge degree”, “number of juvenile felony offenses”, “juvenile misdemeanor offenses”, “other juvenile offenses”, “age”, “sex” and “race” of the defendant. The dataset also contains information on whether the defendant recidivated or not.