



Concentration of Submodular Functions and Read- k Families Under Negative Dependence

Sharmila Duppala  



University of Maryland, College Park, MD, USA

George Z. Li  

Carnegie Mellon University, Pittsburgh, PA, USA

Juan Luque 

University of Maryland, College Park, MD, USA

Aravind Srinivasan  

University of Maryland, College Park, MD, USA

Renata Valieva  

University of Maryland, College Park, MD, USA

Abstract

We study the question of whether submodular functions of random variables satisfying various notions of negative dependence satisfy Chernoff-like concentration inequalities. We prove such a concentration inequality for the lower tail when the random variables satisfy negative association or negative regression, partially resolving an open problem raised in ([23]). Previous work showed such concentration results for random variables that come from specific dependent-rounding algorithms ([6, 14]). We discuss some applications of our results to combinatorial optimization and beyond. We also show applications to the concentration of read- k families [13] under certain forms of negative dependence; we further show a simplified proof of the entropy-method approach of [13].

2012 ACM Subject Classification Mathematics of computing → Probabilistic algorithms

Keywords and phrases Chernoff bounds, Submodular Functions, Negative Correlation

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.47

Funding All authors are supported in part by NSF award number CCF-1918749.

Acknowledgements We thank the ITCS 2025 reviewers for their thoughtful comments.

1 Introduction

Concentration inequalities are ubiquitous in discrete mathematics and theoretical computer science [2, 9]. The most canonical examples are the Chernoff-Hoeffding bounds, which show strong concentration for linear combinations of independent random variables [7, 15]. In some applications, the condition of independence is too restrictive, so weaker notions have been considered [3, 24, 25]. Of interest to us is the setting where the random variables are negatively correlated, which arises naturally, for example, in designing approximation algorithms by solving a linear or semidefinite program and applying some dependent randomized rounding algorithm [11]. For this setting, [20] showed that the Chernoff-Hoeffding bounds can be shown under the weak notion of *negative cylinder dependence*: this and other standard notions of negative dependence are defined in Section 2.1.

For some applications in combinatorial optimization, algorithmic game theory, and machine learning, one needs to consider the more general class of *submodular functions* f of the random variables, rather than simple linear combinations. When the binary random variables X_1, \dots, X_n are independent, it was shown that $f(X_1, \dots, X_n)$ still satisfies Chernoff



© Sharmila Duppala, George Z. Li, Juan Luque, Aravind Srinivasan, and Renata Valieva; licensed under Creative Commons License CC-BY 4.0

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).

Editor: Raghu Meka; Article No. 47; pp. 47:1–47:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

bounds exactly [6]. When there is dependence between the random variables, the results are much weaker. The only known results are for random variables that are output by specific dependent-rounding algorithms, known as swap rounding and pipage rounding [6, 14]. These results showed that a Chernoff-like lower-tail bound also holds for submodular functions for their specific dependent rounding procedure. As noted in the work of [12], it is not clear how to generalize either of these proofs to any general notion of negative dependence.

We introduce a new notion of negative dependence, called 1-negative association, which is weaker than negative association and negative regression but stronger than negative cylinder dependence.

► **Definition 1.** *A collection of random variables X_1, \dots, X_n is said to satisfy 1-negative association if for any two monotone functions f and g , where g depends on a single random variable X_i and f depends on the remaining random variables $\{X_j\}_{j \in [n] \setminus \{i\}}$, we have $\mathbb{E}[fg] \leq \mathbb{E}[f]\mathbb{E}[g]$.*

Importantly, while in general it is weaker than the notion of weak negative regression introduced by [23], 1-negative association is equivalent to it when the variables X_1, \dots, X_n are binary. Further details are provided in Section 3.1.

Our main result is that the Chernoff-like bound shown in [6, 14] also hold under 1-negative association (see Section 3.2). In particular, this implies the following:

► **Theorem 2.** *Let X_1, \dots, X_n be binary random variables with mean x_1, \dots, x_n satisfying negative association (or negative regression). Let f be a non-negative monotone submodular function with marginal values in $[0, 1]$ and let F be the multilinear extension of f . If we let $\mu_0 = F(x_1, \dots, x_n)$, then we have the following:*

$$\Pr[f(X_1, \dots, X_n) \leq (1 - \delta) \cdot \mu_0] \leq \exp(-\mu_0 \delta^2 / 2).$$

A few remarks are in order. First, we highlight that the concentration in the above theorem is with respect to the value of the multilinear extension $F(x_1, \dots, x_n)$, rather than the true expected value $\mathbb{E}[f(X_1, \dots, X_n)]$. In general, the true expected value can be greater than the value of multilinear extension [23]. Nevertheless, this suffices for applications relating to submodular maximization, and is the same type of concentration result shown in previous work. Second, recall that negative cylinder dependence does not suffice to show this concentration bound [6, p. 583]. As a result, our results are, in some informal sense, almost tight in terms of the condition on negative dependence.

In addition to providing submodular concentration results for a wide class of rounding algorithms and distributions, our results also give a new path toward understanding why pipage rounding and swap rounding satisfy the lower-tail Chernoff bound. By proving that the rounding algorithms output random variables which are 1-negatively associated, we immediately obtain a new proof of the lower tail bounds. This can be viewed as evidence that the two rounding algorithms satisfy 1-negative association or even negative association/regression. We leave this as an interesting open question.

Techniques

We use the standard method of bounding the exponential moments for lower-tail Chernoff bounds. Our idea is to show that the exponential moments for our negatively-correlated random variables is upper bounded by the exponential moments for independent copies of the random variables. Formally, let X_1, \dots, X_n be random variables satisfying 1-negative association and let X_1^*, \dots, X_n^* be independent copies of the random variables. We show for any $\lambda < 0$, we have

$$\mathbb{E}[\exp(\lambda \cdot f(X_1, \dots, X_n))] \leq \mathbb{E}[\exp(\lambda \cdot f(X_1^*, \dots, X_n^*))].$$

Since the exponential-moments method has been used to prove Chernoff bounds for submodular functions in the independent case [6], we can then repeat their proof and conclude with our desired result. We believe this proof idea may be of independent interest. For example, the same ideas can show that for a supermodular function g and any $\lambda > 0$, we have

$$\mathbb{E}[\exp(\lambda \cdot g(X_1, \dots, X_n))] \leq \mathbb{E}[\exp(\lambda \cdot g(X_1^*, \dots, X_n^*))].$$

In other words, we have morally proven the following statement: any upper-tail concentration bound which can be proven for a supermodular function g under independence based on the exponential-moments method also holds when the underlying random variables are negatively associated. As an example, we can apply this to a read- k family of supermodular functions g_1, \dots, g_r for negatively associated random variables [13]. A read- k family is defined as a set of functions where each variable appears in at most k functions. This concept is particularly useful in scenarios where functions need to model or manage overlapping sets of variables with constraints on their interaction. We highlight that the proof of concentration for read- k families given in [13] doesn't use the exponential-moments method, but instead it is based on the entropy method. We address this by giving a simpler proof of their results, this time using the exponential moments method. This gives the first concentration results for a class of supermodular functions under negatively correlated random variables, and is detailed in Section 3.3.

Applications

Our motivation for studying the problem comes from the randomized-rounding paradigm in approximation algorithms for converting a fractional solution to a linear program into an integral one. In many such randomized-rounding schemes, the output random variables have been shown to satisfy strong negative dependence properties, such as negative association [26, 11]. For all such rounding algorithms, our results immediately imply the submodular Chernoff lower-tail bound. It remains an interesting open question to efficiently sample negatively dependent distributions for a wider class of set systems. A particularly interesting algorithm is given in the work of [22]; they show that a fractional point in a matroid polytope can be rounded to an integral one such that the resulting distribution preserves marginals and satisfies negative association. However, a gap identified in their proof [23] complicates the application of their approach. The implications of this issue for the applicability of our results remain an area for further investigation.

As a concrete application, we consider the maximum coverage problem under group fairness constraints. Here, we have a universe of elements $\{1, \dots, n\}$, a collection S_1, \dots, S_m of subsets of the universe, and a budget k . We are further given subsets $C_1, \dots, C_\ell \subseteq [n]$ (which should be thought of as demographic groups) along with thresholds w_1, \dots, w_ℓ . Our goal is to choose k sets from the collection to maximize the number of elements covered subject to the fairness constraint each demographic group is sufficiently covered (i.e., at least w_j elements from C_j are covered). Since this is a special case of multiobjective submodular maximization, there exists a $(1 - 1/e - \epsilon)$ -approximation to the problem such that each fairness constraint is approximately satisfied [6, 28]. Unfortunately, these results rely on the randomized swap-rounding algorithm due to its submodular concentration properties, which requires a super-linear time complexity. While swap rounding can be implemented with poly-logarithmic depth [5], a simpler dependent-rounding algorithm of [26] requires linear work and only $O(\log n)$ depth, which improves the efficiency. Observe that the pre-processing step in [28] only requires $O(n\ell)$ time. Since we can solve the linear program for fair maximum

coverage in near-linear time [1], we obtain a near-linear time algorithm for the problem after using the efficient rounding algorithm of [26]. These same ideas can be used to improve the time complexity of the algorithm by [27] for influence maximization with group-fairness constraints. Since the proofs are similar to previous work, we defer the details to a future version of the paper.

More generally, negatively-associated random variables show up naturally in many settings (see e.g., the primer by [29]). [8] studied the canonical example of *balls and bins*, and showed that it satisfied both negative association and negative regression. Another example satisfying the negative-association conditions are any product measure over the set of bases of a balanced matroid, as shown by [10]. A final setting where such random variables occur are random spanning trees, which have been vital in the recent improvements to approximation algorithms for the traveling salesperson problem (see, e.g., [16]). Random spanning trees are known to be strongly Rayleigh, which immediately implies that they are negatively associated. Our results may be interesting here as well.

We also observe that the online rounding scheme of [18] has the strongly Rayleigh property: we immediately get strong concentration (on the lower-tail side) for monotone submodular functions, when the inputs for the function arrive online along with their (Bernoulli) distributions as in the setup of [18].

Related Work

The concentration of negatively-dependent random variables was first formally studied by [19], which showed a central limit theorem for a certain notion of negative dependence. Later on, [20] showed that cylinder negatively dependent random variables yield the Chernoff-Hoeffding concentration inequalities, just like independent random variables. In the context of our paper, these results are somewhat specialized since they focus on linear combinations of random variables.

For non-linear functions of the random variables, the majority of work has focused on the concentration of Lipschitz functions under various notions of negative dependence. [21] showed that for strong Rayleigh measures, one has Gaussian concentration for any Lipschitz function. Later on, [12] corrected an earlier proof of [8], showing that McDiarmid-like concentration results hold for Lipschitz functions of random variables satisfying negative regression. These results are complementary to ours since we are trying to give dimension-free concentration results.

2 Preliminaries

2.1 Notions of Negative Dependence

We begin by defining the notion of negative dependence commonly found in the literature.

Negative Cylinder Dependence

A collection of Boolean random variables X_1, \dots, X_n is said to be negative cylinder dependent if for every $S \subseteq [n]$,

$$\mathbb{E} \left[\prod_{i \in S} X_i \right] \leq \prod_{i \in S} \mathbb{E} [X_i]$$

and

$$\mathbb{E} \left[\prod_{i \in S} (1 - X_i) \right] \leq \prod_{i \in S} \mathbb{E} [1 - X_i].$$

Negative cylinder dependence is the weaker notion considered here. It is known to imply Chernoff bounds for linear combinations of X_1, \dots, X_n but it is insufficient to show our submodular concentration results.

Negative Association

A collection of random variables X_1, \dots, X_n is said to be negatively associated if for any $I, J \subset [n], I \cap J = \emptyset$ and any pair of non-decreasing functions $f : \mathbb{R}^I \rightarrow \mathbb{R}, g : \mathbb{R}^J \rightarrow \mathbb{R}$,

$$\mathbb{E}[f(X_I)g(X_J)] \leq \mathbb{E}[f(X_I)]\mathbb{E}[g(X_J)].$$

Here and in the following, X_S refers to those random variables that are indexed by the elements in S , $X_S = \{X_i : i \in S\}$. Negative association is a significant strengthening of negative cylinder dependence, and has many additional useful closure properties. This will be one of the focuses of the paper.

Negative Regression

A collection of random variables X_1, \dots, X_n is said to satisfy negative regression, if for any $I, J \subset [n], I \cap J = \emptyset$, any non-decreasing function $f : \mathbb{R}^I \rightarrow \mathbb{R}$ and $a \leq b \in \mathbb{R}^J$,

$$\mathbb{E}[f(X_I) \mid X_J = a] \geq \mathbb{E}[f(X_I) \mid X_J = b].$$

Negative regression is a strengthening of negative cylinder dependence, but its relationship with negative association is not yet well understood. It is known that negative association doesn't imply negative regression [8], but the opposite implication is not known. This will be the other focus of the paper.

Strong Rayleigh

A collection of random variables X_1, \dots, X_n is said to satisfy the strong Rayleigh property if the generating function

$$F(z_1, \dots, z_n) = \mathbb{E}[\prod_{j=1}^n z_j^{X_j}]$$

is a real stable polynomial (i.e., it has no root $(z_1, \dots, z_n) \in \mathbb{C}^n$ with all positive imaginary components). The strong Rayleigh property is the strongest notion of negative dependence, and has been shown to imply all other studied negative dependence definitions [4]. As a result, all of our results apply here as well.

2.2 Submodular Functions

We also give a quick review of the basics of submodular functions.

Submodular Functions

We say that a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ is submodular if

$$f(X_1, \dots, X_{i-1}, 1, X_{i+1}, \dots, X_n) - f(X_1, \dots, X_{i-1}, 0, X_{i+1}, \dots, X_n)$$

is a non-increasing function of $X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n$ for each $i \in [n]$. When viewing the binary input of f as the indicator vector for a set, this is equivalent to the more common definition that f is submodular if for any $X, Y \subseteq [n]$ with $X \subseteq Y$ and any $x \notin Y$, we have

$$f(X \cup \{x\}) - f(X) \geq f(Y \cup \{x\}) - f(Y).$$

Supermodular Functions

We say that a function $g : \{0, 1\}^n \rightarrow \mathbb{R}$ is supermodular if

$$g(X_1, \dots, X_{i-1}, 1, X_{i+1}, \dots, X_n) - g(X_1, \dots, X_{i-1}, 0, X_{i+1}, \dots, X_n)$$

is a non-decreasing function of $X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n$ for each $i \in [n]$. When viewing the binary input of g as the indicator vector for a set, this is equivalent to the more common definition that g is supermodular if for any $X, Y \subseteq [n]$ with $X \subseteq Y$ and any $x \notin Y$, we have

$$g(X \cup \{x\}) - g(X) \leq g(Y \cup \{x\}) - g(Y).$$

Multilinear Extension

The multilinear extension of a function f is

$$F(x) = \mathbb{E}[f(x)] = \sum_{S \subseteq [n]} f(S) \prod_{i \in S} x_i \prod_{i \notin S} (1 - x_i),$$

for $x \in [0, 1]^n$. If we view x as a probability vector, the multilinear extension F is simply the expected value of f when each coordinate is rounded independently in $\{0, 1\}$.

3 Submodular Chernoff Bounds

3.1 1-Negative Association and Weak Negative Regression

We first define the weaker notion of negative dependence which we work with, called 1-negative association, and prove some simple properties about it. We also define a related notion of weak negative regression, which is the analogue of 1-negative association for the notion of negative regression, and we show the equivalence between the two for binary random variables and show that weak negative regression is strictly stronger in general. After an initial draft, we discovered that [23] had already introduced the notion of weak negative regression for binary random variables in a context complementary to ours. Using their work, we can immediately show nice properties about 1-negative association.

► **Definition 3.** A collection of random variables X_1, \dots, X_n is said to satisfy 1-negative association if for any two monotone functions f and g , where g depends on a single random variable X_i and f depends on the remaining random variables $\{X_j\}_{j \in [n] \setminus \{i\}}$, we have $\mathbb{E}[fg] \leq \mathbb{E}[f]\mathbb{E}[g]$.

► **Definition 4.** A collection of random variables X_1, \dots, X_n is said to satisfy weak negative regression if for any index i and any monotone function f depending on the remaining random variables $\{X_j\}_{j \in [n] \setminus \{i\}}$, we have $\mathbb{E}[f|X_i = b] \leq \mathbb{E}[f|X_i = a]$ for all $a \leq b$.

In the following lemmata, we show that weak negative regression implies 1-negative association in general. We then show that the reverse implication holds for binary random variables, but give an example showing that it does not hold in general.

▷ **Claim 5.** If a collection of random variables X_1, \dots, X_n satisfies weak negative regression, then it satisfies 1-negative association.

Proof. Assume X_1, \dots, X_n satisfy weak negative regression; we will prove that it also satisfies 1-negative association. Let f and g be monotone functions such that f depends on X_I for some subset $I \subseteq [n]$ and g depends on X_i for $i \notin I$. Without loss of generality, let us assume that f and g are non-decreasing.

First, define the function $F(a) = \mathbb{E}[f(X)|X_i = a]$. Evaluating the expectation $\mathbb{E}[f(X_I)g(X_i)]$, we can express it via the law of total expectation as

$$\mathbb{E}[f(X_I)g(X_i)] = \mathbb{E}[\mathbb{E}[f(X_I)|X_i]g(X_i)] = \mathbb{E}[F(X_i)g(X_i)].$$

Next, we observe that by the definition of weak negative regression, $F(a)$ is non-increasing. Therefore, random variables $F(X_i)$ and $g(X_i)$ are negatively correlated, yielding

$$\mathbb{E}[F(X_i)g(X_i)] \leq \mathbb{E}[F(X_i)]\mathbb{E}[g(X_i)].$$

Converting $\mathbb{E}[F(X_i)]$ back into the terms of f and g , we find

$$\mathbb{E}[F(X_i)]\mathbb{E}[g(X_i)] = \mathbb{E}[f(X)]\mathbb{E}[g(X_i)].$$

The overall inequality $\mathbb{E}[f(X_I)g(X_i)] \leq \mathbb{E}[f(X_I)]\mathbb{E}[g(X_i)]$ thus holds, which establishes that X_1, \dots, X_n are 1-negatively associated, completing the proof. \triangleleft

\triangleright **Claim 6.** If a collection of binary random variables X_1, \dots, X_n satisfies 1-negative association, then it satisfies weak negative regression.

Proof. Recall that we wish to prove that for any non-decreasing function f depending on some subset $I \subseteq [n]$ and any $i \notin I$, we have that

$$\mathbb{E}[f(X_I)|X_i = 0] \geq \mathbb{E}[f(X_I)|X_i = 1].$$

By the definition of 1-negative association, we obtain that for any monotone functions g which depends on X_i , the following inequality holds:

$$\mathbb{E}[f(X_I)g(X_i)] \leq \mathbb{E}[f(X_I)] \cdot \mathbb{E}[g(X_i)]. \quad (1)$$

Without loss of generality, we may assume that $\mathbb{E}[f(X_I)|X_i = 1] = 0$ by shifting f by a constant. By choosing g to be the identity function, we can apply the law of total probability to obtain

$$\mathbb{E}[f(X_I) \cdot g(X_i)] = \Pr[X_i = 0] \cdot \mathbb{E}[f(X_I)|X_i = 0] \cdot g(0) + \Pr[X_i = 1] \cdot \mathbb{E}[f(X_I)|X_i = 1] \cdot g(1) = 0.$$

Plugging this into Equation 1, we obtain

$$0 \leq \mathbb{E}[f(X_I)] \cdot \mathbb{E}[g(X_i)] = \mathbb{E}[f(X_I)|X_i = 0] \Pr[X_i = 0] \cdot \mathbb{E}[X_i],$$

again by the law of total probability. Since $\mathbb{E}[X_i] > 0$ and $\Pr[X_i = 0] > 0$, this implies that

$$\mathbb{E}[f(X_I)|X_i = 0] \geq 0,$$

which concludes the proof since $\mathbb{E}[f(X_I)|X_i = 1] = 0$. \triangleleft

\triangleright **Claim 7.** There exists (non-binary) distributions over 2 random variables which satisfy 1-negative association but not weak negative regression. In other words, 1-negative association is strictly more general than weak negative regression for non-binary random variables.

Proof. Let's first discuss the intuition for the construction of the counterexample. One can show via algebra that $\mathbb{E}[f(X_I)g(X_i)] - \mathbb{E}[f(X_I)]\mathbb{E}[g(X_i)]$ can be expanded as the following expression:

$$\sum_{x < y} \Pr[X_i = x] \Pr[X_i = y] \cdot (\mathbb{E}[f(X_I)|X_i = x] - \mathbb{E}[f(X_I)|X_i = y]) (g(x) - g(y)),$$

where the summation is over the values that X_i takes with non-zero probability.

47:8 Submodular Chernoff Bounds

Now, suppose that for some values $x < y$ we have that

$$\mathbb{E}[f(X_1, \dots, X_n)|X_i = x] - \mathbb{E}[f(X_1, \dots, X_n)|X_i = y] < 0,$$

violating the weak negative regression property. This would imply that some of the summands are positive (since $g(x) < g(y)$ by monotonicity). In the case of binary random variables, the summation would only consist of a single summand so 1-negative association would be violated. For general random variables, the summation consists of multiple terms so the summation may still be negative even when a single summand is positive. Consequently, the random variables may still satisfy 1-negative association.

We now give the example. Consider the random variables (X_1, X_2) which are uniformly distributed on their support set $\{(0, 3), (1, 1), (2, 2), (3, 0)\}$. By considering an identity function $\mathbb{1}_x : \{0, 1, 2, 3\} \rightarrow \{0, 1, 2, 3\}$, we can show that that the distribution of (X_1, X_2) does not satisfy weak negative regression:

$$\mathbb{E}[\mathbb{1}_x(X_2)|X_1 = 1] = \mathbb{1}_x(1) = 1 < 2 = \mathbb{1}_x(2) = \mathbb{E}[\mathbb{1}_x(X_2)|X_1 = 2].$$

However, for any pair of non-decreasing functions $f, g : \{0, 1, 2, 3\} \rightarrow \mathbb{R}$, we have

$$\mathbb{E}[f(X_1)]\mathbb{E}[g(X_2)] - \mathbb{E}[f(X_1)g(X_2)] = \frac{f(1) + f(2) + f(3)}{4} \cdot \frac{g(1) + g(2) + g(3)}{4} - \frac{f(1)g(1) + f(2)g(2)}{4},$$

where we have again assumed without loss of generality that $f(0) = g(0) = 0$.

We claim that the quantity on the right hand side is always non-negative. In order to see this, observe that $f(2)g(2) \leq f(i)g(j)$ for any $i, j \geq 2$ by monotonicity. As a result, we have

$$\frac{f(2)g(2)}{4} \leq \frac{(f(2) + f(3))(g(2) + g(3))}{16}.$$

Further, we observe that $f(1)g(1) \leq f(i)g(j)$ for any $i, j \geq 1$ by monotonicity. As a result, we have

$$\frac{f(1)g(1)}{4} \leq \frac{f(1)(g(2) + g(3)) + g(1)(f(2) + f(3))}{16}.$$

Combining these two inequalities immediately and observing that $f(1)g(1) \geq 0$ by monotonicity implies our desired result. Hence, the distribution is 1-negatively associated. \triangleleft

Since 1-negative association and weak negative regression are equivalent for binary random variables and weak negative regression has been shown to be strictly stronger than cylinder negative dependence [23, Proposition 2.4], we also have that 1-negative association is strictly stronger than cylinder negative dependence. Additionally, since weak negative regression is strictly stronger than 1-negative association for general random variables and weak negative regression has been shown to be strictly weaker than negative association and negative regression [23, Proposition 2.4], we have that 1-negative association is strictly weaker than negative association and negative regression. We summarize these in the following corollaries.

- **Corollary 8.** *1-negative association is a strictly weaker condition than negative association.*
- **Corollary 9.** *1-negative association is a strictly weaker condition than negative regression.*
- **Corollary 10.** *1-negative association is a strictly stronger condition than negative cylinder dependence.*

3.2 Proof of Submodular Concentration

We will now prove our main result. As mentioned in the introduction, our proof is based on the standard technique of bounding the exponential moments. The following lemma contains our main technical contribution, stating that the exponential moments of $f(X_1, \dots, X_n)$ under 1-negative association is dominated by that under independence. Our results will follow easily afterwards.

► **Lemma 11.** *Let X_1, \dots, X_n be 1-negatively associated random variables and let X_1^*, \dots, X_n^* be independent random variables with the same marginal distributions. Also let f be a non-negative monotone function.*

- *If f is submodular and $\lambda < 0$, we have*

$$\mathbb{E}[\exp(\lambda f(X_1, \dots, X_n))] \leq \mathbb{E}[\exp(\lambda f(X_1^*, \dots, X_n^*))].$$
- *If f is supermodular and $\lambda > 0$, we have*

$$\mathbb{E}[\exp(\lambda f(X_1, \dots, X_n))] \leq \mathbb{E}[\exp(\lambda f(X_1^*, \dots, X_n^*))].$$

Proof. Fix $\lambda < 0$ if f is submodular and $\lambda > 0$ if f is supermodular. Observe that in order to prove the lemma, it suffices to prove

$$\mathbb{E}[\exp(\lambda \cdot f(X_1, \dots, X_i, \dots, X_m))] \leq \mathbb{E}[\exp(\lambda \cdot f(X_1, \dots, X_i^*, \dots, X_m))], \quad (2)$$

since we can iteratively apply the above inequality to each X_i (note that we can do this because independent variables are also negatively associated). For simplicity of notation and without loss of generality, we will prove the inequality for $i = 1$.

By considering the cases of $X_1 = 0$ and $X_1 = 1$ separately, we have

$$\exp(\lambda \cdot f(X_1, \dots, X_n)) = X_1 \cdot \exp(\lambda \cdot f(1, X_2, \dots, X_n)) + (1 - X_1) \cdot \exp(\lambda \cdot f(0, X_2, \dots, X_n)),$$

where the equality holds pointwise on the underlying probability space. Via simple algebraic manipulations, we can rewrite the right hand side as

$$X_1 \cdot [\exp(\lambda \cdot f(1, X_2, \dots, X_n)) - \exp(\lambda \cdot f(0, X_2, \dots, X_n))] + \exp(\lambda \cdot f(0, X_2, \dots, X_n))$$

Taking expectations, we now have that $\mathbb{E}[\exp(\lambda f(X_1, \dots, X_n))]$ can be written as

$$\mathbb{E}[X_1 \cdot [\exp(\lambda \cdot f(1, X_2, \dots, X_n)) - \exp(\lambda \cdot f(0, X_2, \dots, X_n))] + \exp(\lambda \cdot f(0, X_2, \dots, X_n))] \quad (3)$$

Observe that X_1 is clearly an increasing function of X_1 . We claim that if either (i) f is submodular and $\lambda < 0$ or (ii) f is supermodular and $\lambda > 0$, we have that $\exp(\lambda \cdot f(1, X_2, \dots, X_n)) - \exp(\lambda \cdot f(0, X_2, \dots, X_n))$ is an increasing function in X_2, \dots, X_n . Indeed, we first rewrite the function as

$$\exp(\lambda \cdot f(0, X_2, \dots, X_n)) \cdot [\exp(\lambda \cdot (f(1, X_2, \dots, X_n) - f(0, X_2, \dots, X_n))) - 1] =: A_1 \cdot A_2$$

for simplicity of notation.

Let us first consider the case when $\lambda < 0$ and f is submodular. We have that A_1 is (i) positive because the exponential function is always positive and (ii) non-increasing in X_2, \dots, X_n because f is non-decreasing and $\lambda < 0$. We also have that A_2 is (i) negative because the argument in $\exp(\cdot)$ is negative, so the exponential is in $(0, 1)$ (ii) non-decreasing since $\lambda < 0$ and the difference of f evaluated at $X_1 = 1$ and $X_1 = 0$ is non-increasing by definition of submodularity. Hence, our expression of interest is the product of a function A_1 which decreases towards 0 and a function A_2 which increases towards 0. The product will be negative and monotonically increasing towards 0.

47:10 Submodular Chernoff Bounds

Now, let us consider the case when $\lambda > 0$ and f is supermodular. We have that A_1 is (i) positive because the exponential function is always positive and (ii) non-decreasing since $\lambda > 0$, f is monotone, and $\exp(\cdot)$ is also monotone. We also have that A_2 is (i) positive because the argument of $\exp(\cdot)$ is positive since f is monotone so the exponential is greater than 1 and (ii) non-decreasing since $\lambda > 0$ and the difference of f evaluated at $X_1 = 1$ and $X_1 = 0$ is non-decreasing by definition of supermodularity. As a result, the product will be positive and non-decreasing, as desired.

Since we have shown that the $A_1 A_2$ is also monotone, we now have that the first term in Equation 3 can be written as the product of monotone functions of disjoint subsets, one of which is the singleton set. By 1-negative association, we have that the first term is upper bounded by

$$\mathbb{E}[X_1] \cdot \mathbb{E}[\exp(\lambda \cdot f(1, X_2, \dots, X_n)) - \exp(\lambda \cdot f(0, X_2, \dots, X_n))].$$

Consequently, the entire expression in (3) is upper bounded by

$$\mathbb{E}[X_1] \cdot \mathbb{E}[\exp(\lambda \cdot f(1, X_2, \dots, X_n)) - \exp(\lambda \cdot f(0, X_2, \dots, X_n))] + \mathbb{E}[\exp(\lambda \cdot f(0, X_2, \dots, X_n))].$$

Since X_1 and X_1^* have the same marginal distributions, the above is exactly equal to

$$\mathbb{E}[X_1^*] \cdot \mathbb{E}[\exp(\lambda \cdot f(1, X_2, \dots, X_n)) - \exp(\lambda \cdot f(0, X_2, \dots, X_n))] + \mathbb{E}[\exp(\lambda \cdot f(0, X_2, \dots, X_n))].$$

And since X_1^* is independent with X_2, \dots, X_m by assumption, the above is equal to

$$\mathbb{E}[X_1^* \cdot \exp(\lambda \cdot f(1, X_2, \dots, X_n)) - \exp(\lambda \cdot f(0, X_2, \dots, X_n))] + \mathbb{E}[\exp(\lambda \cdot f(0, X_2, \dots, X_n))].$$

In particular, observe that this is in the exact same form as Equation 3, except with X_1 replaced with X_1^* . Note that when we transformed the left-hand side of Equation 2 to Equation 3, we never used any properties of the random variables X_1, \dots, X_n other than the fact that they take values in $\{0, 1\}$. As a result, we can reverse the direction of all of the equalities to show that the above expression is equal to

$$\mathbb{E}[\exp(\lambda \cdot f(X_1^*, X_2, \dots, X_n))],$$

which completes the proof of the lemma. \blacktriangleleft

Now, we will complete the proof of our main result. Combining the theorem below with Claims 8 and 9 immediately gives a proof of Theorem 2. Here, our proof will rely heavily on the proof of the Chernoff bound for submodular functions under independence given in [6].

► **Theorem 12.** *Let X_1, \dots, X_n be binary random variables with mean x_1, \dots, x_n satisfying 1-negative association. Let f be a non-negative monotone submodular function with marginal values in $[0, 1]$ and let F be the multilinear extension of f . If we let, $\mu_0 = F(x_1, \dots, x_n)$, then we have the following:*

$$\Pr[f(X_1, \dots, X_n) \leq (1 - \delta) \cdot \mu_0] \leq \exp(-\mu_0 \delta^2 / 2).$$

Proof. Let X_1^*, \dots, X_n^* be independent random variables with the same respective marginals as X_1, \dots, X_n and let $\lambda < 0$ be a parameter to be set later. Let us decompose $f(X_1^*, \dots, X_n^*) = \sum_{i=1}^n Y_i^*$, where

$$Y_i^* = f(X_1^*, \dots, X_i^*, 0, \dots, 0) - f(X_1^*, \dots, X_{i-1}^*, 0, \dots, 0).$$

Let us denote $\mathbb{E}[Y_i^*] = \omega_i$ and $\mu_0 = \sum_{i=1}^n \omega_i = \mathbb{E}[f(X_1^*, \dots, X_n^*)]$. By the convexity of the exponential and the fact that $Y_i^* \in [0, 1]$, we have that

$$\mathbb{E}[\exp(\lambda \cdot Y_i^*)] \leq \omega_i \cdot \exp(\lambda) + (1 - \omega_i) = 1 + [\exp(\lambda) - 1] \cdot \omega_i \leq \exp[(\exp(\lambda) - 1) \cdot \omega_i].$$

Combining the above with Lemma C.1 from [6], we have that

$$\begin{aligned} \mathbb{E}[\exp(\lambda \cdot f(X_1^*, \dots, X_n^*))] &= \mathbb{E}[\exp(\lambda \cdot \sum_{i=1}^n Y_i^*)] \leq \prod_{i=1}^n \mathbb{E}[\exp(\lambda \cdot Y_i^*)] \\ &\leq \exp[(\exp(\lambda) - 1) \cdot \mu_0]. \end{aligned} \quad (4)$$

Now, we can follow the proof of the standard Chernoff bound:

$$\begin{aligned} \Pr[f(X_1, \dots, X_n) \leq (1 - \delta) \cdot \mu_0] &= \Pr[\exp(\lambda \cdot f(X_1, \dots, X_n)) \leq \exp(\lambda(1 - \delta) \cdot \mu_0)] \\ &\leq \frac{\mathbb{E}[\exp(\lambda \cdot f(X_1, \dots, X_n))]}{\exp(\lambda(1 - \delta) \cdot \mu_0)} \\ &\leq \frac{\mathbb{E}[\exp(\lambda \cdot f(X_1^*, \dots, X_n^*))]}{\exp(\lambda(1 - \delta) \cdot \mu_0)} \\ &\leq \frac{\exp[(\exp(\lambda) - 1) \cdot \mu_0]}{\exp(\lambda(1 - \delta) \cdot \mu_0)} \end{aligned}$$

The first equality follows since $\exp(\lambda \cdot x)$ is a monotone function, the first inequality follows by Markov's inequality, the second inequality follows by Lemma 8, and the final inequality follows Equation 4.

Finally, we can choose λ such that $\exp(\lambda) = 1 - \delta$, which gives

$$\Pr[f(X_1, \dots, X_n) \geq (1 - \delta)\mu_0] \leq \frac{\exp(-\delta\mu_0)}{(1 - \delta)^{(1 - \delta)\mu_0}} \leq \exp(-\mu_0 \cdot \delta^2/2),$$

where we used $(1 - \delta)^{1 - \delta} \leq \exp(-\delta + \delta^2/2)$ for $\delta \in (0, 1]$ in the final inequality. \blacktriangleleft

3.3 Concentration of read- k families

In this subsection, we illustrate an application of our proof technique to give concentration for a read- k family of supermodular functions. Read- k families arise naturally in problems such as subgraph counting in random graphs, and can be seen as a complementary weak dependence notion to that of low-degree polynomials [17]. Our work gives the first concentration results for these problems under negative dependence.

Let's consider this notion of weak dependence defined in [13]. Let Y_1, \dots, Y_n be random variables and assume that they can be factored as functions of random variables X_1, \dots, X_m . We say that Y_1, \dots, Y_n are a read- k family of X_1, \dots, X_m if for each variable X_i , there are at most k variables among Y_1, \dots, Y_n that are influenced by X_i . Formally, we have the following.

► **Definition 13.** *Let X_1, \dots, X_m be random variables. For each $j \in [n]$, let $P_j \subseteq [m]$ and let $f_j : \{0, 1\}^{P_j} \rightarrow [0, 1]$ be functions of X_{P_j} . We say that $Y_j = f_j(X_{P_j})$ are a read- k family if $|\{j : i \in P_j\}| \leq k$ for each $i \in [m]$ (i.e., each variable X_i influences at most k functions).*

When X_1, \dots, X_m are independent, [13] showed that we have

$$\Pr[\sum_{j=1}^n f_j(X_{P_j}) \geq (p + \epsilon)n] \leq \exp(-D(p + \epsilon|p) \cdot n/k) \quad (5)$$

$$\Pr[\sum_{j=1}^n f_j(X_{P_j}) \leq (p - \epsilon)n] \leq \exp(-D(p - \epsilon|p) \cdot n/k), \quad (6)$$

47:12 Submodular Chernoff Bounds

where $p = (1/n) \sum_{j=1}^n \mathbb{E}[Y_j]$ and $D(\cdot||\cdot)$ is the Kullback-Leibler divergence. Notably, while Gavinsky et al. do not require random variables X_1, \dots, X_m to be binary, as we do in our approach, their model requires Y_1, \dots, Y_n to be binary. We will show that Inequality 5 on the upper tail still holds for supermodular functions f_1, \dots, f_n . Similarly, Inequality 6, which addresses the lower tail bound, continues to apply to submodular functions f_1, \dots, f_n .

► **Theorem 14.** *Let X_1, \dots, X_m be 1-negatively associated random variables and let X_1^*, \dots, X_m^* be independent random variables with the same respective marginal distributions as X_1, \dots, X_m . Suppose that $f_j(X_{P_j})$ for $j \in [n]$ are a read- k family, where $f_j : \{0, 1\}^{P_j} \rightarrow [0, 1]$ are supermodular functions. If we let $p_0 = (1/n) \sum_{j=1}^n \mathbb{E}[f_j(X_{P_j}^*)]$ denote the averaged expectation when the underlying random variables are independent, we have*

$$\Pr[\sum_{j=1}^n f_j(X_{P_j}) \geq (p_0 + \epsilon)n] \leq \exp(-D(p_0 + \epsilon||p_0) \cdot n/k).$$

Proof. Let $f(X_1, \dots, X_m) = \sum_{j=1}^n f_j(X_{P_j})$ be the quantity of interest, and note that f is the sum of supermodular functions so it is supermodular as well.

We will follow the standard proof via exponential moments. Let $\lambda > 0$; we have

$$\Pr[f(X_1, \dots, X_m) \geq (p_0 + \epsilon)n] = \Pr[\exp(\lambda \cdot f(X_1, \dots, X_m)) \geq \exp(\lambda \cdot (p_0 + \epsilon)n)] \quad (7)$$

$$\leq \mathbb{E}[\exp(\lambda \cdot f(X_1, \dots, X_m))] / \exp(\lambda \cdot (p_0 + \epsilon)n), \quad (8)$$

where the inequality follows by Markov's. Since f is supermodular, we have by Lemma 11 that

$$\mathbb{E}[\exp(\lambda \cdot f(X_1, \dots, X_m))] \leq \mathbb{E}[\exp(\lambda \cdot f(X_1^*, \dots, X_m^*))] = \mathbb{E}[\exp(\lambda \cdot \sum_{j=1}^n f_j(X_{P_j}^*))]. \quad (9)$$

In the new proof of concentration of read- k families, given in the Appendix, we show that

$$\mathbb{E}[\exp(\lambda \cdot \sum_{j=1}^n f_j(X_{P_j}^*))] \leq \left(\prod_{j=1}^n \mathbb{E}[\exp(\lambda \cdot f_j(X_{P_j}^*))^k] \right)^{1/k}. \quad (10)$$

Combining equations 7–10, we have

$$\Pr[f(X_1, \dots, X_m) \geq (p_0 + \epsilon)n] \leq \left(\prod_{j=1}^n \mathbb{E}[\exp(k\lambda \cdot f_j(X_{P_j}^*)) / \exp(k\lambda(p_0 + \epsilon)n)] \right)^{1/k}.$$

Let $\lambda' = k\lambda$; since $\lambda > 0$ is a parameter we set, we can view $\lambda' > 0$ as a parameter as well. We will abuse notation and replace λ' with λ , so we have

$$\Pr[f(X_1, \dots, X_m) \geq (p_0 + \epsilon)n] \leq \left(\prod_{j=1}^n \mathbb{E}[\exp(\lambda \cdot f_j(X_{P_j}^*)) / \exp(\lambda(p_0 + \epsilon)n)] \right)^{1/k},$$

for any $\lambda > 0$. Now, observe that the right hand side of the inequality is the exact same as in the proof of the standard Chernoff bound under independence, except with an additional exponent $1/k$. As a result, we can follow the original proof of the Chernoff bound to show that

$$\Pr[\sum_{j=1}^n f_j(X_{P_j}) \geq (p_0 + \epsilon)n] \leq \exp(-D(p_0 + \epsilon||p_0) \cdot n/k),$$

which was our desired result. ◀

► **Corollary 15.** *Let X_1, \dots, X_m be 1-negatively associated random variables. Suppose that $f_j(X_{P_j})$ for $j \in [n]$ are a read- k family, where $f_j : \{0, 1\}^{P_j} \rightarrow [0, 1]$ are submodular functions. If we let $p_0 = (1/n) \sum_{j=1}^n \mathbb{E}[f_j(X_{P_j}^*)]$ denote the averaged expectation when the underlying random variables are independent, we have*

$$\Pr[\sum_{j=1}^n f_j(X_{P_j}) \leq (p_0 - \epsilon)n] \leq \exp(-D(p_0 - \epsilon||p_0) \cdot n/k).$$

Proof. Define $g_j := 1 - f_j$, where $g_j : \{0, 1\}^{P_j} \rightarrow [0, 1]$ are supermodular. Then $1 - p_0 = (1/n) \sum_{j=1}^n \mathbb{E}[g_j(X_{P_j}^*)]$. Applying Theorem 14, we have:

$$\Pr[\sum_{j=1}^n g_j(X_{P_j}) \geq (1 - (p_0 - \epsilon))n] \leq (1 - p_0 + \epsilon)n \leq \exp(-D(1 - p_0 + \epsilon || 1 - p_0) \cdot n/k).$$

The property of Kullback–Leibler divergence $D(1 - p || 1 - q) = D(p || q)$ implies

$$\begin{aligned} \Pr[\sum_{j=1}^n f_j(X_{P_j}) \leq (p_0 - \epsilon)n] &= \Pr[\sum_{j=1}^n g_j(X_{P_j}) \geq (1 - (p_0 - \epsilon))n] \\ &\leq \exp(-D(p_0 - \epsilon || p_0) \cdot n/k), \end{aligned}$$

which concludes the proof. ◀

References

- 1 Zeyuan Allen-Zhu and Lorenzo Orecchia. Nearly-linear time positive lp solver with faster convergence rate. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC '15*, pages 229–236. Association for Computing Machinery, 2015. doi:10.1145/2746539.2746573.
- 2 Noga Alon and Joel H. Spencer. *The Probabilistic Method, Third Edition*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 2008.
- 3 Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal*, 19(3):357–367, 1967. doi:10.2748/tmj/1178243286.
- 4 Julius Borcea, Petter Brändén, and Thomas M. Liggett. Negative dependence and the geometry of polynomials. *Journal of the American Mathematical Society*, 22(2):521–567, 2009. URL: <http://www.jstor.org/stable/40587241>.
- 5 Chandra Chekuri and Kent Quanrud. Submodular function maximization in parallel via the multilinear relaxation. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 303–322. SIAM, 2019. doi:10.1137/1.9781611975482.20.
- 6 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 575–584. IEEE Computer Society, 2010. doi:10.1109/FOCS.2010.60.
- 7 Herman Chernoff. A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the sum of Observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952. doi:10.1214/aoms/1177729330.
- 8 Devdatt Dubhashi and Desh Ranjan. Balls and bins: a study in negative dependence. *Random Struct. Algorithms*, 13(2):99–124, September 1998. doi:10.1002/(SICI)1098-2418(199809)13:2<99::AID-RSA11%3E3.0.CO;2-M.
- 9 Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. URL: <http://www.cambridge.org/gb/knowledge/isbn/item2327542/>.
- 10 Tomás Feder and Milena Mihail. Balanced matroids. In S. Rao Kosaraju, Mike Fellows, Avi Wigderson, and John A. Ellis, editors, *Proceedings of the 24th Annual ACM Symposium on Theory of Computing, STOC 1992*, pages 26–38. ACM, 1992. doi:10.1145/129712.129716.
- 11 Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *J. ACM*, 53(3):324–360, 2006. doi:10.1145/1147954.1147956.
- 12 Kevin Garbe and Jan Vondrák. Concentration of lipschitz functions of negatively dependent variables, 2018. arXiv:1804.10084.

- 13 Dmitry Gavinsky, Shachar Lovett, Michael E. Saks, and Srikanth Srinivasan. A tail bound for read- k families of functions. *Random Struct. Algorithms*, 47(1):99–108, 2015. doi:10.1002/rsa.20532.
- 14 Nicholas J. A. Harvey and Neil Olver. Pipage rounding, pessimistic estimators and matrix concentration. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, pages 926–945. SIAM, 2014. doi:10.1137/1.9781611973402.69.
- 15 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. URL: <http://www.jstor.org/stable/2282952>.
- 16 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved approximation algorithm for metric TSP. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 32–45. ACM, 2021. doi:10.1145/3406325.3451009.
- 17 Jeong Han Kim and Van H. Vu. Concentration of multivariate polynomials and its applications. *Comb.*, 20(3):417–434, 2000. doi:10.1007/s004930070014.
- 18 Joseph (Seffi) Naor, Aravind Srinivasan, and David Wajc. Online dependent rounding schemes. *arXiv preprint*, 2024. arXiv:2301.08680.
- 19 Charles M. Newman. Asymptotic independence and limit theorems for positively and negatively dependent random variables. *Lecture Notes-Monograph Series*, 5:127–140, 1984. URL: <http://www.jstor.org/stable/4355492>.
- 20 Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff-hoeffding bounds. *SIAM J. Comput.*, 26(2):350–368, 1997. doi:10.1137/S0097539793250767.
- 21 Robin Pemantle and Yuval Peres. Concentration of lipschitz functionals of determinantal and other strong rayleigh measures. *Comb. Probab. Comput.*, 23(1):140–160, 2014. doi:10.1017/S0963548313000345.
- 22 Yuval Peres, Mohit Singh, and Nisheeth K. Vishnoi. Random walks in polytopes and negative dependence. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPICs*, pages 50:1–50:10. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.ITCS.2017.50.
- 23 Frederick Qiu and Sahil Singla. Submodular dominance and applications. In Amit Chakrabarti and Chaitanya Swamy, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022*, volume 245 of *LIPICs*, pages 44:1–44:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.APPROX/RANDOM.2022.44.
- 24 Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM J. Discret. Math.*, 8(2):223–250, 1995. doi:10.1137/S089548019223872X.
- 25 Maciej Skorski. Tight chernoff-like bounds under limited independence. In Amit Chakrabarti and Chaitanya Swamy, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022*, volume 245 of *LIPICs*, pages 15:1–15:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.APPROX/RANDOM.2022.15.
- 26 Aravind Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001*, pages 588–597. IEEE Computer Society, 2001. doi:10.1109/SFCS.2001.959935.
- 27 Alan Tsang, Bryan Wilder, Eric Rice, Milind Tambe, and Yair Zick. Group-fairness in influence maximization. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, pages 5997–6005. ijcai.org, 2019. doi:10.24963/ijcai.2019/831.

- 28 Rajan Udwani. Multi-objective maximization of monotone submodular functions with cardinality constraint. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Annual Conference on Neural Information Processing Systems, NeurIPS 2018*, pages 9513–9524, 2018. URL: <https://proceedings.neurips.cc/paper/2018/hash/7e448ed9dd44e6e22442dac8e21856ae-Abstract.html>.
- 29 David Wajc. Negative association: Definition, properties, and applications. Manuscript, 2017. URL: <https://www.cs.cmu.edu/~ldwajc/notes/Negative%20Association.pdf>.

A Concentration of Read- k Families

We will give a new simpler proof of the results of Gavinsky et al. [13] using exponential moments for f_j for independent random variables X_1, \dots, X_n . Our proof will use the following lemma

► **Lemma 16.** *For an arbitrary read- k family F_1, \dots, F_n , we have*

$$\mathbb{E} \left[\prod_{j=1}^n F_j \right] \leq \left(\prod_{j=1}^n \mathbb{E} [F_j^k] \right)^{1/k}. \quad (11)$$

Using this lemma, if we define

$$F_j = \exp(\lambda \cdot f_j(X_1, \dots, X_m)),$$

it is easy to see that F_1, \dots, F_n are a read- k family since f_1, \dots, f_n are read- k . We will show later that after applying inequality (11) on F_1, \dots, F_m , we can adapt the standard Chernoff bound proof to our case.

Proof. We will prove inequality (11) via induction on the number of independent variables m . For the base case $m = 1$, observe that there are at most k non-constant functions F_j by definition. If there are fewer than k functions F_j , we can also add identity functions without changing the product. As a result, we can without loss of generality assume that there are exactly k functions F_j . The inequality then follows directly by the Generalized Hölder's Inequality.

Now assume we have proven the statement for m independent variables; we will try to prove it for $m + 1$. Again, let $S = S_1$ denote the set of functions F_j which are influenced by X_1 . We have

$$\begin{aligned} \mathbb{E} \left[\prod_{j=1}^n F_j \right] &= \mathbb{E} \left[\prod_{j \in S} F_j \cdot \prod_{j \notin S} F_j \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\prod_{j \in S} F_j \cdot \prod_{j \notin S} F_j \mid X_2, \dots, X_{m+1} \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\prod_{j \in S} F_j \mid X_2, \dots, X_{m+1} \right] \cdot \prod_{j \notin S} F_j \right] \end{aligned} \quad (12)$$

Here, the first equality is obvious, the second equality follows by the law of total expectation, and the third equality follows since F_j for $j \notin S$ only depends on X_2, \dots, X_{m+1} and is independent of X_1 .

After taking the conditional expectation, observe that $\mathbb{E}[\prod_{j \in S} F_j \mid X_2, \dots, X_{m+1}]$ is a random variable which only depends on X_2, \dots, X_{m+1} . In particular, these form a read- k family over m random variables, so we can apply the inductive hypothesis to claim that

$$\mathbb{E} \left[\prod_{j \in S} F_j \mid X_2, \dots, X_{m+1} \right] \leq \left(\prod_{j \in S} \mathbb{E} [F_j^k \mid X_2, \dots, X_{m+1}] \right)^{1/k}$$

47:16 Submodular Chernoff Bounds

After combining this with equation (12), we have

$$\mathbb{E} \left[\prod_{j=1}^n F_j \right] \leq \mathbb{E} \left[\prod_{j \in S} \mathbb{E} [F_j^k | X_2, \dots, X_{m+1}]^{1/k} \cdot \prod_{j \notin S} F_j \right] =: A$$

For $j \notin S$, let $G_j = F_j$ and for $j \in S$, define

$$G_j = \mathbb{E} [F_j^k | X_2, \dots, X_{m+1}]^{1/k}$$

so that $A = \mathbb{E} \left[\prod_{j=1}^n G_j \right]$. Observe that G_j are again a read- k family on X_2, \dots, X_{m+1} , so we can again apply the induction hypothesis to claim that

$$A = \mathbb{E} \left[\prod_{j=1}^n G_j \right] \leq \prod_{j=1}^n \mathbb{E} [G_j^k]^{1/k} = \prod_{j=1}^n \mathbb{E} [F_j^k]^{1/k},$$

where the final equality follows directly by definition of G_j . This completes the proof. ◀

Using the lemma, we can follow the standard Chernoff bound techniques to complete the proof. Let $X = \sum_{j=1}^n f_j(X_1, \dots, X_m)$; we can apply Markov's inequality for any $\lambda < 0$ to obtain

$$\Pr[X \leq qn] = \Pr[\exp(\lambda X) \geq \exp(\lambda qn)] \leq \exp(-\lambda qn) \cdot \mathbb{E}[\exp(\lambda X)]$$

As mentioned before, we can take $F_j = \exp(\lambda \cdot f_j(X_1, \dots, X_m))$ and apply inequality (11) to obtain that

$$\exp(\lambda X) = \mathbb{E} \left[\prod_{j=1}^n F_j \right] \leq \prod_{j=1}^n \mathbb{E}[F_j^k]^{1/k}.$$

Combining this inequality with the previous inequality, we now have that

$$\Pr[X \leq qn] \leq \exp(-\lambda qn) \cdot \prod_{j=1}^n \mathbb{E}[F_j^k]^{1/k} = \left(\prod_{j=1}^n \mathbb{E} \left[\frac{F_j}{\exp(\lambda q)} \right]^k \right)^{1/k}.$$

Writing out the definition of F_j , we have

$$\left[\frac{F_j}{\exp(\lambda q)} \right]^k = \exp(\lambda k \cdot f_j(X_1, \dots, X_m)) / \exp(\lambda k q).$$

Define $\lambda' = \lambda k$; since λ was arbitrary, we will abuse notation and let $\lambda = \lambda'$. Combining the two previous inequalities, we have that

$$\Pr[X \leq qn] \leq \left(\prod_{j=1}^n \mathbb{E}[\exp(\lambda \cdot f_j) / \exp(\lambda q)] \right)^{1/k}.$$

Here, the right-hand side is in exactly the same form as in the proof of the Chernoff-Hoeffding theorem under independence, except with an additional exponent $1/k$. As a result, we can follow the Chernoff-Hoeffding proof and take $q = p - \epsilon$ to obtain

$$\Pr[X \leq (p - \epsilon)n] \leq \exp(-D(p - \epsilon || p) \cdot n/k),$$

which was our desired result.