

RANKED SPREADNESS AND SAMPLE-BASED TESTING

GAIA CARENINI

ABSTRACT. In this note, we introduce the notion of *ranked spreadness*, a strengthening of the usual spread condition in which the elements of each member can be ordered so that their one-coordinate marginals decay geometrically with their rank. This additional structure removes the dependence on the maximum set size in random-containment estimates. We prove width-free hitting and weighted-concentration theorems for ranked-spread set systems, together with an elementary kernel-extraction theorem showing that ranked spreadness arises naturally in arbitrary distributions on small sets.

Our main application is to the simulation of nonadaptive property testers by sample-based testers. If a one-sided tester has average query complexity d and rejects every far input with probability at least δ , then, for every integer $c > d/\delta$, it admits a one-sided sample-based simulation with expected sample complexity $O_{d,\delta,|\Sigma|}(n^{1-1/c})$. More generally, if positive inputs are rejected with probability at most γ and far inputs with probability at least $\delta > \gamma$, the same conclusion holds for every $c > d/(\delta - \gamma)$. In particular, for constant-query nonadaptive testers we obtain an exponent $1 - \Theta(1/q)$, matching, up to the dependence on the rejection gap, the exponent conjectured by Fischer, Lachish, and Vasudev.

1. INTRODUCTION

Spreadness is a basic pseudorandomness condition for set systems. In one standard normalization, a weighted family $\mathcal{F}$ is (m, k) -spread if

$$\mu(\{F \in \mathcal{F} : A \subseteq F\}) \leq mk^{-|A|}$$

for every nonempty set A . Such bounds control the clustering of members of $\mathcal{F}$ and underlie the robust-sunflower method, expectation-threshold results, and a number of sampling and local-to-global arguments [2, 11, 5, 10, 9].

A persistent feature of general spread-to-hitting theorems is a dependence on the maximum size of a member. If every member has size at most r , ordinary spreadness typically guarantees random containment only at density of order $(\log r)/k$. This logarithmic loss is necessary for unrestricted spread families, but it can be artificial when the spread condition comes from an underlying geometric hierarchy of one-coordinate marginals.

In this work, we isolate such a hierarchy through the notion of *ranked spreadness*. A weighted indexed family $(P_t)_{t \in \mathcal{T}}$ is ranked-spread if, for every t , the elements of P_t can be ordered as $P_t = \{x_1(t), \dots, x_{|P_t|}(t)\}$ so that the marginal of the element in rank i is at most mk^{-i} . The ordering is allowed to depend on the member P_t ; there need not be a global rank assigned to the ground-set elements. Ranked spreadness implies ordinary spreadness with the same parameters, but retains considerably more information.

The reason why this extra information is useful is fairly simple. When two members intersect, we charge their interaction to the last-ranked element of the first member lying in the second. At sampling density $p = \alpha/k$, this bounds the Janson dependency sum by a multiple of

$$\sum_{i=1}^r (pk)^{-i} = \sum_{i=1}^r \alpha^{-i}.$$

For every $\alpha > 1$, this is bounded independently of r . Thus ranked spreadness converts the width-dependent loss in the usual spread lemma into a convergent geometric series.

Date: August 12, 2026.

1.1. The ranked-spread method. We develop three complementary forms of this principle. The first is a *random-hitting theorem*. If an indexed family is (m, k) -ranked-spread and $\mathcal{A}$ is a subfamily of mass τ , then a Bernoulli sample of density $p = \alpha/k$ contains some member of $\mathcal{A}$ with failure probability at most $\exp(-\tau/2m \sum_{i=1}^r \alpha^{-i})$. In particular, whenever $\alpha > 1$, the failure probability is at most $\exp(-(\alpha - 1)\tau/2m)$, with no dependence on the maximum member size r .

The second is a *weighted-concentration theorem*. For a sampled set W , define

$$X_{\mathcal{A}} = \sum_{t \in \mathcal{A}} \mu(t) p^{-|P_t|} \mathbf{1}_{\{P_t \subseteq W\}}.$$

This is an unbiased estimator for the mass $\mu(\mathcal{A})$. Ranked spreadness controls both its variance and its lower tail, again through the geometric series $\sum_i \alpha^{-i}$. This weighted statement is strictly more informative than the assertion that at least one member is exposed, since it permits one to estimate the total mass of an indexed subfamily from a random sample.

The third ingredient is a *kernel-extraction theorem*. Given an arbitrary distribution on sets of average size d , and an integer $c > d$, we find a scale $k = n^{1/c}$, a kernel K of size at most $n^{1-1/c}$, and a subfamily of mass at least $1 - d/c$ such that the petals obtained after removing K are ranked-spread. The proof is elementary. We divide the one-coordinate marginals into geometric scales and use an interval-packing argument to choose a boundary at which most members exhibit the required rank-by-rank decay.

1.2. Robust daisies. Our first application is a strengthened robust daisy lemma. Robust daisies were introduced by Goldberg, Gur, and Saraogi [6] in order to retain random-hitting guarantees simultaneously for every subfamily of a distribution on small sets. Their framework already yields one-sided sample-based simulations, but its general extraction theorem incurs losses depending on the set width and on the size of the support.

Ranked kernel extraction preserves the more precise marginal hierarchy hidden in the extraction procedure. Combining it with the ranked hitting theorem gives robustness at every density $p = \alpha n^{-1/c}$ with $\alpha > 1$, and produces a failure exponent independent of both the support size and the maximum petal width. In the normalization of [6, Lemma 5.1], their stated bound applies for $\alpha > 2q$ and has exponent proportional to $\alpha/(q^2 \log |\mathcal{D}|) \cdot |K|$. Our bound applies for every $\alpha > 1$ and has exponent proportional to $(\alpha - 1)/q \cdot B$, where B is the extraction scale and may be larger than the kernel itself. Thus the ranked-spread formulation both simplifies the extraction and identifies the parameter naturally governing robustness.

1.3. Application to sample-based testing. Our principal application concerns the conversion of nonadaptive property testers into sample-based testers. Property testing, introduced by Goldreich, Goldwasser, and Ron [7], asks how efficiently one can distinguish inputs satisfying a property from inputs far from it. A standard tester chooses a small set of coordinates to query. A *sample-based tester* instead observes the values on a uniform Bernoulli sample of the coordinates and then performs arbitrary post-processing. This model was systematically studied by Goldreich and Ron [8], who asked when arbitrary testers can be simulated from random samples.

Fischer, Lachish, and Vasudev [4] proved that every constant-query nonadaptive tester over a fixed alphabet has a sublinear sample-based simulation. In the one-sided case, their direct argument uses sampling probability of order n^{-1/q^2} , and they conjectured that the correct exponent should be $1/q$. Dall'Agnol, Gur, and Lachish [3] subsequently treated adaptive robust local algorithms, obtaining exponent $1/O(q^2 \log^2 q)$ in that more general setting. The robust-daisy machinery of Goldberg, Gur, and Saraogi also implies a one-sided conversion, with slightly weaker quantitative parameters than those obtained here.

We prove a linear, rather than quadratic, dependence on the query complexity. More precisely, our bounds depend on the *average* number of queries, as well as on the rejection gap.

Let $\Pi \subseteq \Sigma^n$ be a property over a finite alphabet, and let $r_T(x)$ denote the rejection probability of a nonadaptive tester T on input x . We first state the one-sided result.

Theorem 1.1 (One-sided query-to-sample conversion). *Let Σ be a finite alphabet, let $\varepsilon, \delta \in (0, 1]$, and let T be a one-sided nonadaptive $\varepsilon/2$ -tester for $\Pi \subseteq \Sigma^n$. Suppose that T has average query complexity $d > 0$ and rejects every $\varepsilon/2$ -far input with probability at least δ . Let c be an integer with $c > d/\delta$, and let $\beta = \delta - \frac{d}{c} > 0$, and $\alpha_{\text{hit}} = 1 + \frac{2d}{\beta}(\log |\Sigma| + \log 3)$. Then, for all sufficiently large n , Π has a one-sided ε -tester which samples every coordinate independently with probability $p = \alpha_{\text{hit}} n^{-1/c}$. It rejects every ε -far input with probability at least $2/3$ and has expected sample complexity $\alpha_{\text{hit}} n^{1-1/c}$.*

The one-sided simulation uses only the hitting form of ranked spreadness. After extracting a small kernel, the new tester considers every possible assignment to the kernel. For each false assignment, the rejecting tests form a subfamily of positive mass, and the ranked hitting theorem ensures that the sample exposes one of their petals. Perfect completeness guarantees that the true kernel assignment is never eliminated.

The two-sided setting requires more. A valid input may itself expose rejecting tests, so the existence of a single rejecting local view carries no useful information. Instead, the tester must estimate the mass of rejecting tests associated with each possible kernel assignment. This is precisely the role of weighted ranked concentration.

Theorem 1.2 (Two-sided query-to-sample conversion). *Let Σ be a finite alphabet, let $\varepsilon \in (0, 1]$, and let $0 \leq \gamma < \delta \leq 1$. Let T be a nonadaptive tester for $\Pi \subseteq \Sigma^n$. Suppose that T has average query complexity $d > 0$, rejects every input in Π with probability at most γ , and rejects every $\varepsilon/2$ -far input with probability at least δ . Let c be an integer with $c > \frac{d}{\delta - \gamma}$, and let $\beta = \delta - \frac{d}{c}$, $g = \beta - \gamma > 0$, and $\alpha_{\text{wt}} = 1 + \frac{4d}{g^2}(3\gamma + 2\log(3|\Sigma|))$. Then, for all sufficiently large n , Π has a two-sided ε -tester which samples every coordinate independently with probability $p = \alpha_{\text{wt}} n^{-1/c}$. It accepts every input in Π with probability at least $2/3$, rejects every ε -far input with probability at least $2/3$, and has expected sample complexity $\alpha_{\text{wt}} n^{1-1/c}$.*

Taking $d \leq q$ gives the following convenient worst-case formulation.

Corollary 1.3. *A one-sided q -query tester with detection probability δ admits a one-sided sample-based simulation with expected sample complexity $O_{q,\delta,|\Sigma|}(n^{1-1/(\lfloor q/\delta \rfloor + 1)})$. More generally, a two-sided q -query tester with positive-input rejection probability at most γ and far-input rejection probability at least δ admits a simulation with expected sample complexity $O_{q,\gamma,\delta,|\Sigma|}(n^{1-1/(\lfloor q/(\delta-\gamma) \rfloor + 1)})$. For $\gamma = 1/3$ and $\delta = 2/3$, this is $O_{q,|\Sigma|}(n^{1-1/(3q+1)}) = n^{1-\Theta(1/q)}$.*

The distinction between the two conversions is important. When $\gamma = 0$, the weighted theorem recovers the same exponent as the one-sided theorem, but with a sampling constant proportional to the inverse square of the retained rejection margin. The one-sided argument uses perfect completeness and only needs to expose one witness for each false kernel assignment. The two-sided argument has no such Boolean elimination rule and instead relies on the importance-weighted estimator supplied by ranked concentration. In this sense, the two testing results illustrate the two basic probabilistic consequences of ranked spreadness: hitting and mass estimation.

1.4. Organization. In section 2, we introduce the testing model and recall the unweighted and weighted forms of Janson's inequality used throughout. In section 3, we define ranked spreadness and prove the width-free hitting and weighted-concentration theorems. In section 4, we prove ranked kernel extraction, derive its hitting and concentration consequences, and obtain the improved robust daisy lemma. In section 5, we apply the framework to one-sided and two-sided query-to-sample conversion.

Acknowledgments. The author is grateful to Guy Goldberg, Tom Gur, and Sidhant Saraogi for helpful discussions that clarified the relationship between the present work and theirs. She is especially grateful to Tom Gur for encouraging her to investigate applications of an earlier version of the main result, which led to a substantial strengthening of the paper. This work was supported by the CB European PhD Studentship funded by Trinity College, Cambridge.

2. PRELIMINARIES

Let Σ be a finite alphabet. For $x, y \in \Sigma^n$, their normalized Hamming distance is

$$\text{dist}(x, y) = \frac{1}{n} |\{i \in [n] : x_i \neq y_i\}|.$$

For a property $\Pi \subseteq \Sigma^n$, write

$$\text{dist}(x, \Pi) = \min_{y \in \Pi} \text{dist}(x, y).$$

The word x is ε -far from Π if $\text{dist}(x, \Pi) > \varepsilon$.

A deterministic *local test* is a pair $t = (Q_t, \varphi_t)$, where $Q_t \subseteq [n]$ and $\varphi_t : \Sigma^{Q_t} \rightarrow \{0, 1\}$. We interpret $\varphi_t = 1$ as rejection. A finite probability distribution μ on a set $\mathcal{T}$ of local tests defines a nonadaptive randomized tester: it samples $t \sim \mu$, queries Q_t , and outputs $\varphi_t(x|_{Q_t})$. Its rejection probability on x is $r_T(x) = \mu(\{t \in \mathcal{T} : \varphi_t(x|_{Q_t}) = 1\})$. The maximum and average query complexities are, respectively, $q = \max_{t \in \text{supp}(\mu)} |Q_t|$ and $d = \mathbb{E}_{t \sim \mu} |Q_t|$.

Every finite-randomness nonadaptive tester has this form after including all random coins in the index t . More generally, since there are only finitely many pairs consisting of a query set of size at most q and a deterministic predicate on that set, arbitrary internal randomness can be grouped into a finite distribution of local tests.

The tester is *one-sided* for Π if $r_T(x) = 0$ for every $x \in \Pi$. Equivalently, every rejecting local view of every positive-probability local test is inconsistent with membership in Π . For $p \in [0, 1]$, let $W \sim \text{Bin}([n], p)$ denote the random set obtained by retaining each coordinate independently with probability p . A *p-sampling tester* receives $(W, x|_W)$ and applies arbitrary post-processing. Its expected sample complexity is pn .

This is the sample-based model of Goldreich and Ron [8]. We do not impose a running-time bound on the post-processing.

The combinatorial arguments use the query sets but must retain the indexing. In particular, two local tests may use the same query set and different predicates. Thus we work with a probability space $(\mathcal{T}, \mu)$ and a set $P_t \subseteq U$ attached to every index $t \in \mathcal{T}$. For $x \in U$, write $\mu(x) = \mu(\{t \in \mathcal{T} : x \in P_t\})$. For $\mathcal{A} \subseteq \mathcal{T}$, we use $\mu(\mathcal{A})$ for its probability mass.

We use the following standard form of Janson's inequality; see, for example, [1, Chapter 8].

Lemma 2.1 (Janson). *Let $\mathcal{M}$ be a finite family of subsets of a finite ground set, and let W contain each ground-set element independently with probability p . For $M \in \mathcal{M}$, let $Z_M = \mathbf{1}_{\{M \subseteq W\}}$, let $Z = \sum_{M \in \mathcal{M}} Z_M$, $\lambda = \mathbb{E}Z$, and define the ordered dependency sum*

$$\Delta = \sum_{M, N \in \mathcal{M}: M \neq N, M \cap N \neq \emptyset} \mathbb{E}[Z_M Z_N].$$

Then

$$\mathbb{P}[Z = 0] \leq \exp \left(- \min \left\{ \frac{\lambda}{2}, \frac{\lambda^2}{2\Delta} \right\} \right),$$

with the second term interpreted as $+\infty$ when $\Delta = 0$.

Lemma 2.2 (Weighted lower-tail Janson). *Let $(A_i)_{i \in I}$ be a finite indexed family of nonempty subsets of a finite ground set, let W include every ground-set element independently with*

probability p , and let $w_i \geq 0$. Let $I_i = \mathbf{1}_{\{A_i \subseteq W\}}$, $X = \sum_{i \in I} w_i I_i$, $\lambda = \mathbb{E}X$, and define the ordered dependency sum, including the diagonal, by

$$\overline{\Delta} = \sum_{\substack{i, j \in I \\ A_i \cap A_j \neq \emptyset}} w_i w_j \mathbb{E}[I_i I_j].$$

Then, for every $\eta \in [0, 1]$,

$$\mathbb{P}[X \leq (1 - \eta)\lambda] \leq \exp\left(-\frac{\eta^2 \lambda^2}{2\overline{\Delta}}\right), \quad (2.1)$$

with the natural interpretation when $\overline{\Delta} = 0$.

Proof. Suppose first that every w_i is rational. Choose N so that Nw_i is an integer for all i , and replace A_i by Nw_i indexed copies. If Y is the number of sampled copies, then $Y = NX$, and $\mathbb{E}Y = N\lambda$. The ordered dependency sum for the copied family, including the diagonal, is $N^2\overline{\Delta}$. The standard lower-tail Janson inequality therefore gives

$$\mathbb{P}[X \leq (1 - \eta)\lambda] = \mathbb{P}[Y \leq (1 - \eta)\mathbb{E}Y] \leq \exp\left(-\frac{\eta^2 N^2 \lambda^2}{2N^2 \overline{\Delta}}\right).$$

For arbitrary nonnegative weights, approximate the weights by nonnegative rationals. Since the index set is finite, the random variables, their means, and their dependency sums converge uniformly; passing to the limit, with an arbitrarily small slack in the threshold, gives (2.1). $\square$

3. RANKED SPREADNESS AND RANDOM HITTING

We start by introducing the notion of *ranked spreadness*.

Definition 3.1 (Ranked spreadness). Let $m, k > 0$. A weighted indexed set system $(P_t)_{t \in \mathcal{T}}$ with distribution μ is (m, k) -*ranked-spread* if every nonempty P_t admits an ordering $P_t = \{x_1(t), \dots, x_{|P_t|}(t)\}$ such that, for every i , $\mu(\{s \in \mathcal{T} : x_i(t) \in P_s\}) \leq mk^{-i}$.

In the normalization used here, a weighted indexed set system is (m, k) -*spread* if $\mu(\{s \in \mathcal{T} : A \subseteq P_s\}) \leq mk^{-|A|}$ for every nonempty $A \subseteq U$. This is the usual spread condition from the robust-sunflower and expectation-threshold literature [2, 11, 5].

Ranked spreadness implies ordinary spreadness with the same parameters. To see this, suppose that $A \subseteq P_t$, and let $x_i(t)$ be the element of A having largest rank in the ranked-spread ordering of P_t . Since A contains $|A|$ elements, $i \geq |A|$, and hence

$$\mu(\{s : A \subseteq P_s\}) \leq \mu(\{s : x_i(t) \in P_s\}) \leq mk^{-i} \leq mk^{-|A|}.$$

Thus ranked spreadness is a genuine strengthening of ordinary spreadness.

For an arbitrary spread system whose members have size at most r , the general spread lemma gives random containment at density of order $(\log r)/k$; equivalently, the general expectation-threshold theory retains a logarithmic dependence on the largest minimal witness [5, 10, 9]. The additional edgewise hierarchy is what allows theorem 3.2 to work at every density α/k with $\alpha > 1$, with an exponential failure bound independent of r . This hierarchy is distinct from other multiscale refinements of spreadness, such as Spiro's smoother spread condition [12], which controls codegrees across prescribed ranges of intersection sizes.

The ordering may depend on t ; there need not be a globally consistent rank assigned to each coordinate. For $r \geq 1$ and $\alpha > 0$, let $H_r(\alpha) = \sum_{i=1}^r \alpha^{-i}$.

Theorem 3.2 (Ranked spread-to-hitting). Let $(P_t)_{t \in \mathcal{T}}$ be an (m, k) -ranked-spread indexed set system, and suppose that $|P_t| \leq r$ for every t . Let $\mathcal{A} \subseteq \mathcal{T}$, put $\tau = \mu(\mathcal{A})$, and let $W \sim \text{Bin}(U, p)$, and $p = \frac{\alpha}{k} \leq 1$. Then

$$\mathbb{P}_W[\nexists t \in \mathcal{A} : P_t \subseteq W] \leq \exp\left(-\frac{\tau}{2mH_r(\alpha)}\right). \quad (3.1)$$

If $\alpha > 1$, then

$$\mathbb{P}_W[\#t \in \mathcal{A} : P_t \subseteq W] \leq \exp\left(-\frac{(\alpha-1)\tau}{2m}\right). \quad (3.2)$$

Proof. If $\tau = 0$, there is nothing to prove. If $P_t = \emptyset$ for some $t \in \mathcal{A}$, the event on the left has probability zero. We may therefore assume that every set indexed by $\mathcal{A}$ is nonempty.

Condition on $\mathcal{A}$. The resulting distribution $\mu_{\mathcal{A}}$ is $(m/\tau, k)$ -ranked-spread, using the same edgewise orderings. It is therefore enough to prove the theorem when $\mathcal{A} = \mathcal{T}$ and then replace m by m/τ .

We first suppose that all values $\mu(t)$ are rational. Choose a multiset $\mathcal{M}_0$ in which the proportion of copies of each index t is $\mu(t)$. Let $r_0 = \max_{t \in \mathcal{T}} |P_t| \leq r$.

For an integer $L \geq 1$, replace every member of $\mathcal{M}_0$ by L copies. Pad each copy of P_t to size r_0 by adding $r_0 - |P_t|$ fresh dummy elements, used only by that copy. Let $\mathcal{M}_L$ be the resulting uniform multiset on the enlarged ground set. The dummy elements are mutually distinct and do not belong to U .

Sample every element of the enlarged ground set independently with probability p . If a padded set is sampled, then its original set is contained in $W \cap U$. Hence

$$\mathbb{P}_W[\#t : P_t \subseteq W] \leq \mathbb{P}[\text{no member of } \mathcal{M}_L \text{ is sampled}]. \quad (3.3)$$

Let $N_L = |\mathcal{M}_L|$ and, for $M \in \mathcal{M}_L$, let Z_M be the indicator that M is sampled. Then

$$\lambda = \mathbb{E} \sum_{M \in \mathcal{M}_L} Z_M = N_L p^{r_0}.$$

Fix $M \in \mathcal{M}_L$, and order its original elements according to the ranked-spread ordering of the corresponding P_t : $x_1(M), \dots, x_{|P_t|}(M)$. If another padded member N intersects M , their intersection contains no dummy element. Let $i(M, N) = \max\{i : x_i(M) \in N\}$. Then $|M \cap N| \leq i(M, N)$. Moreover, the number of $N \in \mathcal{M}_L$ with $i(M, N) = i$ is at most the number containing $x_i(M)$, and hence at most $N_L m k^{-i}$. Consequently, the ordered Janson dependency sum satisfies

$$\Delta \leq \sum_{M \in \mathcal{M}_L} \sum_{i=1}^{r_0} N_L m k^{-i} p^{2r_0-i} = N_L^2 p^{2r_0} m \sum_{i=1}^{r_0} (pk)^{-i} = \lambda^2 m H_{r_0}(\alpha) \leq \lambda^2 m H_r(\alpha). \quad (3.4)$$

By theorem 2.1,

$$\mathbb{P}[Z = 0] \leq \exp\left(-\min\left\{\frac{N_L p^{r_0}}{2}, \frac{1}{2m H_r(\alpha)}\right\}\right).$$

The first term tends to infinity as $L \rightarrow \infty$. Combining this with (3.3) gives

$$\mathbb{P}_W[\#t : P_t \subseteq W] \leq \exp\left(-\frac{1}{2m H_r(\alpha)}\right).$$

For a general finite distribution, approximate its atom weights by positive rational weights with the same support. The coordinate marginals converge, so the approximating systems are $(m + o(1), k)$ -ranked-spread. The event to be bounded depends only on the indexed support and the sets P_t , not on the weights. Letting the approximation error tend to zero proves (3.1).

Finally, if $\alpha > 1$, then

$$H_r(\alpha) = \sum_{i=1}^r \alpha^{-i} \leq \frac{1}{\alpha - 1},$$

which gives (3.2). $\square$

Observe that, at the critical density $p = 1/k$, the theorem gives an exponent of order $1/(mr)$. Further, for $p = \alpha/k$ with $\alpha > 1$, the dependence on r disappears.

3.1. Weighted ranked concentration. The hitting theorem only records whether some petal is exposed. For two-sided testing we need to estimate the total mass of a rejecting subfamily. The following importance-weighted version provides exactly this lower-tail control.

Theorem 3.3 (Weighted ranked concentration). *Let $(P_t)_{t \in \mathcal{T}}$ be an (m, k) -ranked-spread indexed set system with distribution μ , and suppose that $|P_t| \leq r$ for every t . Let $\mathcal{A} \subseteq \mathcal{T}$, put $\tau = \mu(\mathcal{A})$, and let $W \sim \text{Bin}(U, p)$, with $p = \frac{\alpha}{k} \leq 1$. Define $X_{\mathcal{A}} = \sum_{t \in \mathcal{A}} \mu(t) p^{-|P_t|} \mathbf{1}_{\{P_t \subseteq W\}}$. Then, $\mathbb{E}X_{\mathcal{A}} = \tau$, $\text{Var}(X_{\mathcal{A}}) \leq \tau m H_r(\alpha)$, and, for every $\eta \in [0, 1]$,*

$$\mathbb{P}[X_{\mathcal{A}} \leq (1 - \eta)\tau] \leq \exp\left(-\frac{\eta^2 \tau}{2m H_r(\alpha)}\right). \quad (3.5)$$

Proof. Indices with $P_t = \emptyset$ contribute deterministically to $X_{\mathcal{A}}$. Let their total mass be τ_0 , and put $\tau_+ = \tau - \tau_0$. If $(1 - \eta)\tau < \tau_0$, then the event in (3.5) is empty. Otherwise it is a lower-tail event for the positive-petal contribution with relative deviation $\eta_+ = \frac{\eta\tau}{\tau_+} \leq 1$. Since $\eta_+^2 \tau_+ = \eta^2 \tau^2 / \tau_+ \geq \eta^2 \tau$, it is enough to treat the case in which every P_t , $t \in \mathcal{A}$, is nonempty.

For $t \in \mathcal{A}$, set $I_t = \mathbf{1}_{\{P_t \subseteq W\}}$, and $w_t = \mu(t) p^{-|P_t|}$. Then $X_{\mathcal{A}} = \sum_{t \in \mathcal{A}} w_t I_t$, and

$$\mathbb{E}[w_t I_t] = \mu(t) p^{-|P_t|} p^{|P_t|} = \mu(t),$$

which proves that $\mathbb{E}X_{\mathcal{A}} = \tau$.

Define

$$\Lambda_{\mathcal{A}} := \sum_{\substack{t, s \in \mathcal{A} \\ P_t \cap P_s \neq \emptyset}} w_t w_s \mathbb{E}[I_t I_s] = \sum_{\substack{t, s \in \mathcal{A} \\ P_t \cap P_s \neq \emptyset}} \mu(t) \mu(s) p^{-|P_t \cap P_s|}. \quad (3.6)$$

Fix $t \in \mathcal{A}$, and use a ranked-spread ordering $P_t = \{x_1(t), \dots, x_{|P_t|}(t)\}$. For every s with $P_s \cap P_t \neq \emptyset$, let $i(t, s) = \max\{i : x_i(t) \in P_s\}$. Then $|P_t \cap P_s| \leq i(t, s)$. Since $p \leq 1$, $p^{-|P_t \cap P_s|} \leq p^{-i(t, s)}$. Moreover,

$$\sum_{\substack{s \in \mathcal{A} \\ i(t, s) = i}} \mu(s) \leq \mu(\{s : x_i(t) \in P_s\}) \leq m k^{-i}.$$

Consequently,

$$\sum_{\substack{s \in \mathcal{A} \\ P_t \cap P_s \neq \emptyset}} \mu(s) p^{-|P_t \cap P_s|} \leq m \sum_{i=1}^{|P_t|} k^{-i} p^{-i} \leq m H_r(\alpha). \quad (3.7)$$

Multiplying by $\mu(t)$ and summing over $t \in \mathcal{A}$ gives

$$\Lambda_{\mathcal{A}} \leq \tau m H_r(\alpha). \quad (3.8)$$

If $P_t \cap P_s = \emptyset$, then I_t and I_s are independent. For an intersecting pair, $\text{Cov}(I_t, I_s) \leq \mathbb{E}[I_t I_s]$. It follows from (3.8) that $\text{Var}(X_{\mathcal{A}}) \leq \Lambda_{\mathcal{A}} \leq \tau m H_r(\alpha)$, proving $\text{Var}(X_{\mathcal{A}}) \leq \tau m H_r(\alpha)$.

Applying theorem 2.2 with the weights w_t and using (3.8), we obtain

$$\mathbb{P}[X_{\mathcal{A}} \leq (1 - \eta)\tau] \leq \exp\left(-\frac{\eta^2 \tau^2}{2\Lambda_{\mathcal{A}}}\right) \leq \exp\left(-\frac{\eta^2 \tau}{2m H_r(\alpha)}\right).$$

□

4. RANKED KERNEL EXTRACTION

We now show that every distribution on small indexed sets contains a large ranked-spread subfamily after the removal of a small kernel. The boundary selection is inspired by the robust-daisy extraction of Goldberg, Gur, and Saraogi [6]. The new point is to retain the ordered one-coordinate marginal information produced by the proof.

Lemma 4.1 (Interval packing). *Let A be a finite multiset of integers. Call an integer j bad if there is $r \geq 1$ such that $|A \cap \{j - r, j - r + 1, \dots, j - 1\}| \geq r$, where multiplicities are counted. Then the number of bad integers is at most $|A|$.*

Proof. Greedily process the bad integers from right to left. Let j be the largest remaining bad integer, and choose $r \geq 1$ witnessing its badness. Remove all remaining bad integers in $\{j - r + 1, \dots, j\}$. At most r bad integers are removed, while the interval $\{j - r, \dots, j - 1\}$ contains at least r elements of A , counted with multiplicity. Charge the removed bad integers to distinct such elements.

At the next step, the chosen bad integer is at most $j - r$, so its charging interval lies strictly to the left of the preceding one. Thus the charging intervals are pairwise disjoint. The total number of removed bad integers is therefore at most $|A|$. $\square$

Theorem 4.2 (Ranked kernel extraction). *Let $(Q_t)_{t \in \mathcal{T}}$ be an indexed set system on a ground set U of size n , equipped with a probability distribution μ . Suppose that $|Q_t| \leq q$ for every t , and let $d = \mathbb{E}_{t \sim \mu} |Q_t|$. Assume $d > 0$, let $c > d$ be an integer, and set $k = n^{1/c}$. Then there exist $j \in \{1, \dots, c\}$, a subfamily $D \subseteq \mathcal{T}$, a kernel $K \subseteq U$, and a scale $B = n/k^j = n^{1-j/c}$ such that $\mu(D) \geq 1 - d/c$, and $|K| \leq B \leq n^{1-1/c}$. Moreover, the indexed petal system $P_t = Q_t \setminus K$, for all $t \in D$, under the conditioned distribution μ_D , is $(d/\mu(D)B, k)$ -ranked-spread.*

Proof. For $x \in U$, define its one-coordinate marginal $\delta(x) = \mu(\{t : x \in Q_t\})$. Then

$$\sum_{x \in U} \delta(x) = \mathbb{E}_{t \sim \mu} |Q_t| = d. \quad (4.1)$$

For every integer s , let $\lambda_s = d \frac{k^s}{n}$. For $j \in \{1, \dots, c\}$, define $K_j = \{x \in U : \delta(x) \geq \lambda_j\}$. By (4.1), $|K_j| \lambda_j \leq d$, and hence

$$|K_j| \leq \frac{d}{\lambda_j} = \frac{n}{k^j}. \quad (4.2)$$

We say that j is *good* for an index t if every nonempty $R \subseteq Q_t \setminus K_j$ contains some $x \in R$ satisfying

$$\delta(x) < \lambda_{j-|R|}. \quad (4.3)$$

Otherwise, j is bad for t .

Fix $t \in \mathcal{T}$. Every $x \in Q_t$ has positive marginal, because t has positive probability. Choose an integer $\ell(x)$ such that $\lambda_{\ell(x)} \leq \delta(x) < \lambda_{\ell(x)+1}$, and let A_t be the multiset of the levels $\ell(x)$, for $x \in Q_t$. If j is bad for t , there is a nonempty $R \subseteq Q_t \setminus K_j$, say $|R| = r$, such that $\delta(x) \geq \lambda_{j-r}$ for every $x \in R$. Since $R \cap K_j = \emptyset$, we also have $\delta(x) < \lambda_j$. Therefore $j - r \leq \ell(x) \leq j - 1$ for every $x \in R$. The interval $\{j - r, \dots, j - 1\}$ thus contains at least r elements of A_t . By theorem 4.1, the number of integers j that are bad for t is at most $|A_t| = |Q_t|$.

Let $D_j = \{t \in \mathcal{T} : j \text{ is good for } t\}$. Averaging over $j \in [c]$ and then over $t \sim \mu$ gives

$$\begin{aligned} \frac{1}{c} \sum_{j=1}^c \mu(D_j) &= 1 - \frac{1}{c} \mathbb{E}_{t \sim \mu} |\{j \in [c] : j \text{ is bad for } t\}| \\ &\geq 1 - \frac{1}{c} \mathbb{E}_{t \sim \mu} |Q_t| \\ &= 1 - \frac{d}{c}. \end{aligned}$$

Choose j with $\mu(D_j) \geq 1 - d/c$, and let $D = D_j$, $K = K_j$, and $B = n/k^j$. The asserted kernel bounds follow from (4.2).

It remains to verify ranked spreadness. Fix $t \in D$, and order the elements of its petal $P_t = Q_t \setminus K$ by nonincreasing marginal: $\delta(x_1(t)) \geq \delta(x_2(t)) \geq \dots \geq \delta(x_{|P_t|}(t))$.

For $i \in [|P_t|]$, apply (4.3) to $R_i = \{x_1(t), \dots, x_i(t)\}$. Some $x \in R_i$ satisfies $\delta(x) < \lambda_{j-i}$. Since $x_i(t)$ has the smallest marginal in R_i ,

$$\delta(x_i(t)) < \lambda_{j-i} = d \frac{k^{j-i}}{n} = \frac{d}{B} k^{-i}.$$

As $x_i(t) \notin K$, membership in a petal is the same as membership in the original query set. Therefore

$$\mu_D(\{s \in D : x_i(t) \in P_s\}) \leq \frac{\delta(x_i(t))}{\mu(D)} < \frac{d}{\mu(D)B} k^{-i}.$$

This is precisely the claimed ranked-spread property. $\square$

Combining extraction with theorem 3.2 gives the form used in the testing argument.

Corollary 4.3 (Ranked kernel hitting). *Under the hypotheses and notation of theorem 4.2, let $\mathcal{A} \subseteq D$ and let $W \sim \text{Bin}(U, \frac{\alpha}{k})$, $\alpha/k \leq 1$. If $r = \max_{t \in D} |Q_t \setminus K|$, then*

$$\mathbb{P}_W[\nexists t \in \mathcal{A} : Q_t \setminus K \subseteq W] \leq \exp\left(-\frac{B\mu(\mathcal{A})}{2dH_r(\alpha)}\right). \quad (4.4)$$

In particular, if $\alpha > 1$, then

$$\mathbb{P}_W[\nexists t \in \mathcal{A} : Q_t \setminus K \subseteq W] \leq \exp\left(-\frac{(\alpha-1)B}{2d}\mu(\mathcal{A})\right). \quad (4.5)$$

Proof. The petal distribution under μ_D is $(d/(\mu(D)B), k)$ -ranked-spread. The μ_D -mass of $\mathcal{A}$ is $\mu(\mathcal{A})/\mu(D)$. Substitution into theorem 3.2 cancels the factor $\mu(D)$ and gives (4.4). The final assertion follows from $H_r(\alpha) \leq 1/(\alpha-1)$. $\square$

4.1. Weighted kernel concentration. The weighted analogue of theorem 4.3 follows by applying the preceding theorem under the conditioned petal distribution.

Corollary 4.4 (Weighted ranked kernel concentration). *Under the notation of theorem 4.2, let $\mathcal{A} \subseteq D$, write $\tau = \mu(\mathcal{A})$, and let $W \sim \text{Bin}(U, \frac{\alpha}{k})$, and $\frac{\alpha}{k} \leq 1$. Define*

$$X_{\mathcal{A}} = \sum_{t \in \mathcal{A}} \mu(t) \left(\frac{\alpha}{k}\right)^{-|P_t|} \mathbf{1}_{\{P_t \subseteq W\}}.$$

Then

$$\mathbb{E}X_{\mathcal{A}} = \tau, \quad \text{Var}(X_{\mathcal{A}}) \leq \frac{\tau d}{B} H_q(\alpha),$$

and, for every $\eta \in [0, 1]$,

$$\mathbb{P}[X_{\mathcal{A}} \leq (1-\eta)\tau] \leq \exp\left(-\frac{\eta^2 \tau B}{2dH_q(\alpha)}\right). \quad (4.6)$$

Proof. Apply theorem 3.3 under μ_D . Its weighted score is $X_{\mathcal{A}}/\mu(D)$, its subfamily mass is $\tau/\mu(D)$, and its ranked parameter is $d/(\mu(D)B)$. The factors of $\mu(D)$ cancel in both the variance and lower-tail estimates. $\square$

4.2. A robust daisy corollary. We now translate the preceding result into the terminology of Goldberg, Gur, and Saraogi [6]. The definition is stated for distributions on sets rather than indexed set systems, since this is the form used in their work.

Definition 4.5 (Robust daisy). Let ν be a probability distribution on a finite family of subsets of a ground set U , let $p, \eta \in (0, 1)$, and let $K \subseteq U$. We say that ν is a (p, η) -robust daisy with kernel K if, for every $\mathcal{A} \subseteq \text{supp}(\nu)$,

$$\mathbb{P}_{W \sim \text{Bin}(U, p)}[\exists S \in \mathcal{A} : S \setminus K \subseteq W] \geq 1 - \eta^{\nu(\mathcal{A})}. \quad (4.7)$$

Equivalently, the event in (4.7) is that $S \subseteq K \cup W$ for some $S \in \mathcal{A}$. The sets $S \setminus K$ are called the *petals*.

Thus robustness is required not only for the full support, but for every subfamily, with a failure probability that decays exponentially in its probability mass.

Corollary 4.6 (Improved robust daisy lemma). *Let μ be a probability distribution on subsets of an n -element ground set U , suppose that every set in $\text{supp}(\mu)$ has size at most q , and put $d = \mathbb{E}_{S \sim \mu} |S| > 0$. Let $c > d$ be an integer and set $k = n^{1/c}$. Then there exist $\mathcal{D} \subseteq \text{supp}(\mu)$, a kernel $K \subseteq U$, an index $j \in \{1, \dots, c\}$, and $B = \frac{n}{k^j} = n^{1-j/c}$ such that*

$$\mu(\mathcal{D}) \geq 1 - \frac{d}{c}, \quad |K| \leq B \leq n^{1-1/c}. \quad (4.8)$$

Moreover, for every $\alpha > 1$ for which $p = \frac{\alpha}{k} = \alpha n^{-1/c} \leq 1$, the conditioned distribution $\mu_{\mathcal{D}}$ is a (p, η) -robust daisy with kernel K , where

$$\eta = \exp \left(-\frac{(\alpha - 1)\mu(\mathcal{D})B}{2d} \right). \quad (4.9)$$

In particular, if only the worst-case bound $|S| \leq q$ is retained and $c > q$, then

$$\mu(\mathcal{D}) \geq 1 - \frac{q}{c}, \quad \eta \leq \exp \left(-\frac{(\alpha - 1)(1 - q/c)}{2q} |K| \right). \quad (4.10)$$

Proof. Apply theorem 4.2 to the indexed system whose indices are the sets in $\text{supp}(\mu)$ and whose attached set at index S is S . This gives $\mathcal{D}, K, j$, and B satisfying (4.8). Let $\mathcal{A} \subseteq \mathcal{D}$. By theorem 4.3,

$$\begin{aligned} \mathbb{P}_W[\exists S \in \mathcal{A} : S \setminus K \subseteq W] &\leq \exp \left(-\frac{(\alpha - 1)B}{2d} \mu(\mathcal{A}) \right) \\ &= \left[\exp \left(-\frac{(\alpha - 1)\mu(\mathcal{D})B}{2d} \right) \right]^{\mu_{\mathcal{D}}(\mathcal{A})} \\ &= \eta^{\mu_{\mathcal{D}}(\mathcal{A})}. \end{aligned}$$

This is precisely theorem 4.5. Finally, $d \leq q$, $\mu(\mathcal{D}) \geq 1 - d/c \geq 1 - q/c$, and $B \geq |K|$, which give (4.10). $\square$

The preceding corollary is the promised robust-daisy deduction. In the normalization of Goldberg, Gur, and Saraogi [6, Lemma 5.1], their stated extraction applies for $\alpha > 2q$ and gives the robustness parameter

$$\exp \left(-\frac{\alpha(1 - q/c)}{8q^2 \log |\mathcal{D}|} |K| \right).$$

By contrast, theorem 4.6 applies for every $\alpha > 1$ and gives the support-size-free exponent

$$\frac{(\alpha - 1)(1 - q/c)}{2q} |K|.$$

More strongly, the natural exponent in (4.9) is proportional to the extraction scale B , which may be larger than the kernel itself. The improvement is exactly where ranked spreadness is used.

5. SAMPLE-BASED SIMULATION

The two reductions use the same extracted kernel but exploit it in different ways. We therefore prove them separately.

5.1. The one-sided reduction.

Proof of theorem 1.1. Represent T as a probability distribution μ on indexed deterministic local tests $t = (Q_t, \varphi_t)$. Apply theorem 4.2 with the chosen integer c . We obtain a retained family $D \subseteq \mathcal{T}$, a kernel K , and a scale B such that

$$\mu(D) \geq 1 - \frac{d}{c}, \quad |K| \leq B \leq n^{1-1/c}. \quad (5.1)$$

Let $W \sim \text{Bin}([n], p)$ with $p = \alpha_{\text{hit}} n^{-1/c}$. For sufficiently large n , $p \leq 1$.

Given $(W, x|_W)$, process every hypothetical assignment $a \in \Sigma^K$. Call a *eliminated* if there is a test $t \in D$ such that $P_t := Q_t \setminus K \subseteq W$ and t rejects the local word obtained by combining a on $Q_t \cap K$ with the observed values on P_t . Reject if and only if every assignment is eliminated.

If $x \in \Pi$, the true assignment $a_* = x|_K$ is never eliminated, since every exposed view is an actual view of x and the original tester is one-sided. Thus completeness is perfect.

Now suppose that x is ε -far from Π . For $a \in \Sigma^K$, define $x^{(a)}$ by replacing $x|_K$ with a . Since $|K| \leq n^{1-1/c} \leq \varepsilon n/2$ for all sufficiently large n , every $x^{(a)}$ is $\varepsilon/2$ -far from Π . Hence its rejecting family $\mathcal{R}_a = \{t \in \mathcal{T} : \varphi_t(x^{(a)}|_{Q_t}) = 1\}$ has mass at least δ , and therefore

$$\mu(D \cap \mathcal{R}_a) \geq \delta - \frac{d}{c} = \beta. \quad (5.2)$$

By theorem 4.3,

$$\mathbb{P}[a \text{ is not eliminated}] \leq \exp\left(-\frac{(\alpha_{\text{hit}} - 1)\beta B}{2d}\right).$$

There are at most $|\Sigma|^{|K|} \leq \exp(B \log |\Sigma|)$ assignments. A union bound and the definition of α_{hit} give

$$\mathbb{P}[\text{the new tester accepts } x] \leq \exp\left(B \log |\Sigma| - \frac{(\alpha_{\text{hit}} - 1)\beta B}{2d}\right) = e^{-B \log 3} \leq \frac{1}{3}.$$

Finally, $\mathbb{E}|W| = pn = \alpha_{\text{hit}} n^{1-1/c}$. □

5.2. The two-sided reduction. The Boolean elimination rule cannot be used when $\gamma > 0$: the true kernel assignment may itself expose rejecting tests. Instead, for each kernel assignment we estimate the retained rejection mass and compare it with a threshold lying strictly between the completeness and soundness regimes.

Proof of theorem 1.2. Represent T as a distribution μ on indexed deterministic local tests $t = (Q_t, \varphi_t)$ and apply theorem 4.2. We obtain D, K, B satisfying

$$\mu(D) \geq 1 - \frac{d}{c}, \quad |K| \leq B \leq n^{1-1/c}. \quad (5.3)$$

Let $p = \alpha_{\text{wt}} n^{-1/c}$, $W \sim \text{Bin}([n], p)$, and $\theta = \beta + \gamma/2$. For sufficiently large n , $p \leq 1$.

For $a \in \Sigma^K$, let $x^{(a)}$ agree with a on K and with x outside K , and let

$$\mathcal{R}_a = \{t \in \mathcal{T} : \varphi_t(x^{(a)}|_{Q_t}) = 1\}.$$

Define the observable weighted rejection score

$$Z_a = \sum_{t \in D \cap \mathcal{R}_a} \mu(t) p^{-|P_t|} \mathbf{1}_{\{P_t \subseteq W\}}, \quad P_t = Q_t \setminus K. \quad (5.4)$$

When $P_t \subseteq W$, the predicate value in the summand is determined by the sample together with a . The tester accepts if and only if

$$\min_{a \in \Sigma^K} Z_a \leq \theta. \quad (5.5)$$

Completeness. Suppose $x \in \Pi$ and take $a_* = x|_K$. Put $\tau_* = \mu(D \cap \mathcal{R}_{a_*})$. Then $\tau_* \leq r_T(x) \leq \gamma$. By theorem 4.4,

$$\mathbb{E}Z_{a_*} = \tau_*, \quad \text{Var}(Z_{a_*}) \leq \frac{\tau_* d}{B} H_q(\alpha_{\text{wt}}) \leq \frac{\gamma d}{B} H_q(\alpha_{\text{wt}}).$$

Since $\theta - \tau_* \geq g/2$, Chebyshev's inequality and $H_q(\alpha) \leq (\alpha - 1)^{-1}$ give

$$\begin{aligned} \mathbb{P}[\text{the new tester rejects } x] &\leq \mathbb{P}[Z_{a_*} > \theta] \\ &\leq \frac{4\gamma d H_q(\alpha_{\text{wt}})}{Bg^2} \\ &\leq \frac{4\gamma d}{(\alpha_{\text{wt}} - 1)g^2} \leq \frac{1}{3}. \end{aligned}$$

The case $\gamma = 0$ is included: then $\tau_* = 0$ and $Z_{a_*} = 0$ almost surely.

Soundness. Suppose x is ε -far from Π . As above, every completion $x^{(a)}$ is $\varepsilon/2$ -far for sufficiently large n . Therefore

$$\tau_a := \mu(D \cap \mathcal{R}_a) \geq \delta - \frac{d}{c} = \beta. \quad (5.6)$$

Since $\theta < \tau_a$, apply theorem 4.4 with relative deviation $1 - \theta/\tau_a$. Using $\tau_a \leq 1$ and $\tau_a - \theta \geq g/2$, we obtain

$$\begin{aligned} \mathbb{P}[Z_a \leq \theta] &\leq \exp\left(-\frac{B(\tau_a - \theta)^2}{2dH_q(\alpha_{\text{wt}})\tau_a}\right) \\ &\leq \exp\left(-\frac{Bg^2}{8dH_q(\alpha_{\text{wt}})}\right) \\ &\leq \exp\left(-\frac{(\alpha_{\text{wt}} - 1)Bg^2}{8d}\right). \end{aligned} \quad (5.7)$$

A union bound over at most $|\Sigma|^{|K|} \leq e^{B \log |\Sigma|}$ assignments and the definition of α_{wt} give

$$\mathbb{P}[\text{the new tester accepts } x] \leq \exp\left(B \log |\Sigma| - \frac{(\alpha_{\text{wt}} - 1)Bg^2}{8d}\right) \leq e^{-B \log 3} \leq \frac{1}{3}.$$

Finally, $\mathbb{E}|W| = pn = \alpha_{\text{wt}} n^{1-1/c}$. $\square$

Proof of theorem 1.3. Use $d \leq q$. In the one-sided case choose $c = \lfloor q/\delta \rfloor + 1$. In the two-sided case choose $c = \lfloor q/(\delta - \gamma) \rfloor + 1$. For $\gamma = 1/3$ and $\delta = 2/3$, the latter gives $c = 3q + 1$. $\square$

Note that setting $\gamma = 0$ in theorem 1.2 gives the same exponent as theorem 1.1, but the weighted proof requires $\alpha_{\text{wt}} - 1 = O(d\beta^{-2} \log |\Sigma|)$, whereas the hitting proof requires only $\alpha_{\text{hit}} - 1 = O(d\beta^{-1} \log |\Sigma|)$. This is not a cosmetic constant loss: it reflects the distinction between merely finding one rejecting witness and estimating rejection mass to additive accuracy comparable with the gap.

REFERENCES

- [1] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. 4th ed. Wiley, 2016.
- [2] Ryan Alweiss, Shachar Lovett, Kewen Wu, and Jiapeng Zhang. “Improved Bounds for the Sunflower Lemma”. In: *Annals of Mathematics* 194.3 (2021), pp. 795–815.
- [3] Marcel Dall’Agnol, Tom Gur, and Oded Lachish. “A Structural Theorem for Local Algorithms with Applications to Coding, Testing, and Verification”. In: *SIAM Journal on Computing* 52.6 (2023), pp. 1413–1463.

- [4] Eldar Fischer, Oded Lachish, and Yadu Vasudev. “Trading Query Complexity for Sample-Based Testing and Multi-testing Scalability”. In: *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2015)*. IEEE, 2015, pp. 1163–1182.
- [5] Keith Frankston, Jeff Kahn, Bhargav Narayanan, and Jinyoung Park. “Thresholds versus Fractional Expectation-Thresholds”. In: *Annals of Mathematics* 194.2 (2021), pp. 475–495.
- [6] Guy Goldberg, Tom Gur, and Sidhant Saraogi. “Nearly Tight Lower Bounds for Relaxed Locally Decodable Codes via Robust Daisies”. In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*. Association for Computing Machinery, 2026, pp. 2210–2217.
- [7] Oded Goldreich, Shafi Goldwasser, and Dana Ron. “Property Testing and Its Connection to Learning and Approximation”. In: *Journal of the ACM* 45.4 (1998), pp. 653–750.
- [8] Oded Goldreich and Dana Ron. “On Sample-Based Testers”. In: *ACM Transactions on Computation Theory* 8.2 (2016), 7:1–7:54.
- [9] Elchanan Mossel, Jonathan Niles-Weed, Nike Sun, and Ilias Zadik. “A Bayesian Proof of the Spread Lemma”. In: *Random Structures & Algorithms* 66.4 (2025), e70008.
- [10] Jinyoung Park and Huy Tuan Pham. “A Proof of the Kahn–Kalai Conjecture”. In: *Journal of the American Mathematical Society* 37.1 (2024), pp. 235–243.
- [11] Anup Rao. “Coding for Sunflowers”. In: *Discrete Analysis* 2 (2020), pp. 1–8.
- [12] Sam Spiro. “A Smoother Notion of Spread Hypergraphs”. In: *Combinatorics, Probability and Computing* 32.5 (2023), pp. 809–818.

TRINITY COLLEGE CAMBRIDGE, DEPARTMENT OF PURE MATHEMATICS AND MATHEMATICAL STATISTICS,
CENTRE FOR MATHEMATICAL SCIENCES, WILBERFORCE ROAD, CAMBRIDGE CB3 0WA, UNITED KINGDOM
Email address: gc645@cam.ac.uk