Toward a KKL Theorem for Any HDX

Max Hopkins
Institute for Advanced Study, Princeton, NJ, USA

Abstract

The KKL Theorem, a seminal result in boolean function analysis, characterizes the structure of low-influence (non-expanding) functions on the hypercube. While recent years have seen breakthrough results across a variety of areas relying on analogs of the KKL Theorem *beyond the cube* (e.g. on product spaces, Grassmann graphs), further progress has been inhibited by our poor understanding of the phenomenon across more general domains. Motivated in this context, Bafna, Hopkins, Kaufman, and Lovett (STOC 2022) and Gur, Lifshitz, and Liu (STOC 2022) proved a generalized KKL-type Theorem for spectral high dimensional expanders (HDX). Their results, however, remain highly restricted due to strong quantitative expansion requirements on the underlying complex.

In this work, we introduce a simple local-to-global method for analyzing low influence functions on simplicial complexes. Using this method we prove a local-to-global KKL-type Theorem: any simplicial complex whose links satisfy a KKL-Theorem also satisfies such a result globally. Building on Gotlib and Kaufman (RANDOM 2023), we also prove a weaker dimension-dependent KKL-type Theorem for simplicial complexes with *any non-trivial (two-sided) expansion*. As concrete applications of our framework, we give the first characterization of non-expanding functions on "combinatorial" HDX such as dense clique complexes and a corresponding Kruskal-Katona Theorem, as well as a small-set expansion theorem for the Ramanujan Complexes of Lubotzky, Samuels, and Vishne (EJC '05).

2012 ACM Subject Classification Theory of computation

Keywords and phrases High Dimensional Expanders, Boolean Function Analysis

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2026.74

Category RANDOM

Related Version *Full Version*: `https://arxiv.org/abs/2606.29449`

Funding *Max Hopkins*: Supported by NSF Award DMS-2424441.

1 Introduction

In 1988, Kahn, Kalai, and Linial [51] proved the KKL Theorem: a characterization of low influence (non-expanding) functions[1] on the hypercube that revolutionized the field of boolean function analysis. In recent years, variants of the KKL Theorem *beyond the cube*, including on products [13, 70, 38, 36, 46, 56, 63], Lie groups [30], and the Grassmann [62, 29] have become a powerful tool across a surprisingly broad array of areas, leading, e.g., to the theory of sharp thresholds [37], multiple breakthroughs in PCP theory [62, 61, 27, 26, 10, 60, 66], new streaming lower bounds [32], and to progress on longstanding problems in group theory and combinatorics [30, 58, 57, 59, 43]. Despite its broad importance our understanding of KKL-type theorems on general spaces remains poor, standing as a barrier to further progress on key problems such as the unique games conjecture [10].

Toward this end, a series of works over the past several years [55, 19, 6, 7, 44, 39, 47] developed a new framework for unifying our understanding of Fourier analysis and KKL-type Theorems beyond the cube: *spectral high dimensional expanders* (HDX). HDX are

[1] Informally, a low-influence function on the hypercube is one which has low expected probability of changing value under a random coordinate flip.

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2026).
Editors: Mohit Singh and Tom Gur; Article No. 74; pp. 74:1–74:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

a generalization of expanders to hypergraphs and ranked posets that have seen a recent explosion of application within theoretical computer science [55, 19, 6, 7, 44, 39, 25, 18, 53, 2, 24, 48, 11, 5, 3, 4, 16, 15, 14, 33, 50, 64, 12, 23, 68, 28, 21, 20, 8, 9, 49, 45, 22]. In particular, Bafna, Hopkins, Kaufman, and Lovett [6, 7], Gur, Lifshitz, and Liu [44], and Hopkins [47] proved a KKL-Type Theorem for certain spectral high dimensional expanders stating that any low influence function must be *local*, highly concentrated in some restriction of the hypergraph. Unfortunately, these results only hold for the regime of *near-perfect* expansion, a strong requirement satisfied only by a few families of sparse hypergraphs [65, 54, 17].

In this work, we develop simple local-to-global machinery for analyzing low influence functions on hypergraphs allowing us to remove these requirements in certain well-studied settings. We prove two main results: a KKL-Type Theorem for simplicial complexes whose links (local restrictions) themselves satisfy such a bound such as quotients of the affine building [65] (the so-called Ramanujan complexes), and a weaker (dimension-dependent) KKL-type Theorem for simplicial complexes with *any* non-trivial (two-sided) spectral high dimensional expansion allowing us to give the first non-trivial characterization of low influence functions on product-based HDX [40] and dense clique complexes. We also derive Kruskal-Katona theorems for these objects, a broadly used set of results in extremal combinatorics bounding the lower shadow of a set system.

Our work builds on a growing body of Fourier analytic tools for high dimensional expanders [55, 19, 6, 7, 44, 39, 18, 2, 42] and relies in particular on a decomposition theorem of Gotlib and Kaufman (GK) [42]. We make two main technical contributions in this line. First, we show for any hypergraph X whose *negative* local spectra is bounded away from -1, "global" functions on X (i.e. functions not concentrated in any local component) have bounded projection onto low levels of the GK-decomposition. Second, we consider an *inductive* approach for bounding expansion through a new "function-dependent" variant of Garland's Lemma, a way to write the expansion of f as an expectation across local parts of the complex where, unlike prior such methods, the underlying distribution may depend on the function f in question. In a similar vein, our local-to-global theorem for (global) small-set expansion combines this latter Garland method with Fourier analytic properties of a localized noise operator to derive global expansion bounds.

2 Background

2.1 High Dimensional Expanders

We focus on weighted *pure simplicial complexes* (X, Π) where

$$X = X(0) \cup \ldots \cup X(d)$$

for $X(d) \subset \binom{n}{d}$ an arbitrary d-uniform hypergraph and $X(i) \subset \binom{n}{i}$ given by downward closure, and

$$\Pi = (\pi_0, \ldots, \pi_d)$$

for π_d an arbitrary distribution over $X(d)$ and π_i given by selecting a uniformly random i-set from $\tau \sim \pi_d$.

Our bounds will depend on the now standard notion of high dimensional expansion which bounds the behavior of local neighborhoods of X called *links*. For every face $\tau \in X(i)$, the link of τ is given by:

$$X_\tau \coloneqq \{\sigma \setminus \tau \in X : \tau \subseteq \sigma \in X\}$$

Given a function $f : X(k) \to \mathbb{R}$, its localization $f|_\tau : X_\tau(k - |\tau|) \to \mathbb{R}$ is $f_\tau(\sigma) = f(\tau \cup \sigma)$.

Denote the normalized adjacency matrix of the graph underlying X_τ by A_τ, and for $0 \leq j \leq d-2$ denote the worst-case spectral parameters of j-links by

$$\gamma_j \coloneqq \max_{\tau \in X(j)} \{\lambda_2(A_\tau)\}, \qquad \gamma_j^{(-)} \coloneqq \min_{\tau \in X(j)} \{\lambda_{\min}(A_\tau)\}.$$

A complex is called *strongly connected* if every $\gamma_j < 1$, a (one-sided) *local-spectral expander* if every γ_j is bounded away from 1 [52, 67, 25], and a (two-sided) local-spectral expander if every $\gamma_j^{(-)}$ is additionally bounded away from -1.

2.2 Influence and the Down-Up Walk

As in the seminal work of KKL, we are interested in characterizing the structure of boolean functions $C_k(X, \mathbb{F}_2) \coloneqq \{f : X(k) \to \{0,1\}\}$ with low total influence. In our setting it will be convenient to use an equivalent definition (up to normalization) concerning the expansion of a celebrated random walk on complexes called the *lower* or *down-up* walk. The lower walk at level k, denoted N_k^1, moves between k-faces of the complex via a shared $(k-1)$-face, removing a vertex uniformly at random and re-sampling from the appropriate conditional distribution on the remaining $(k-1)$-face. The *edge expansion* of a set $S \subset X(k)$ with respect to the lower walk is the expected probability of leaving S in a single step, or equivalently:

$$\Phi(S) = 1 - \frac{\langle N_k^1 1_S, 1_S \rangle}{\langle 1_S, 1_S \rangle}.$$

It is not hard to see that expansion, which can also be written as $\frac{\langle 1_S, L 1_S \rangle}{\langle 1_S, 1_S \rangle}$ for the "Laplacian Operator" L, is equivalent to standard notions of influence up to normalization (see full version for details), so we will focus on this notion throughout instead.

2.3 Small-Set Expansion and the Noise Operator

The small-set expansion theorem is a classical and closely related result to the KKL theorem [1, 51]. In its standard form, it states that small sets on the noisy cube expand near perfectly. To prove an analog result beyond the cube, we need to define a corresponding notion of the noise operator T_ρ on simplicial complexes. A natural generalization is to consider the walk on k-faces of X which fixes each vertex with probability ρ, and re-samples the remaining vertices. Formally, we may express this operator as a convex combination of "longer" down-up walks

$$T_\rho \coloneqq \sum_{i=0}^{k} \binom{k}{i} (1-\rho)^{k-i} \rho^i N_k^{k-i},$$

where N_k^j moves between k-faces via a shared $(k-j)$-face. Note this recovers the standard operator when X is a product space (see [7]). We will always write expansion with respect to the noise operator at Φ_{T_ρ}, and write just Φ for the down-up walk. Finally, it will sometimes be useful to reference the opposite variant of the lower walks, the "upper" walks $\widehat{N}_k^i$, which moves between k-faces via a shared $(k+i)$-face.

3 Results

We are now ready to more formally cover our results. We split the subsection into two parts: structure theorems for low-influence functions on arbitrary HDX, and a local-to-global small-set expansion theorem.

3.1 A Structure Theorem for Low Influence Functions

The canonical examples of low influence (non-expanding) sets on simplicial complexes are *links*.[2] Our goal is to show these are the *only* such examples. To this end, we call a function f (ε, i)-global if it is uncorrelated with i-links:

$$\forall \tau \in X(i) : \mathop{\mathbb{E}}_{X_\tau}[f_\tau] \leq \mathbb{E}[f] + \varepsilon.$$

We prove that global functions indeed expand, where the quantitative parameters depend on the underlying local-spectral expansion of the complex.

▶ **Theorem 1** (Expansion of Global Functions). *Let (X, Π) be a k-dimensional simplicial complex and $f : X(k) \to \mathbb{R}$ any (ε, i)-global boolean function. Then*

$$\Phi(f) \geq \frac{1 - \mathbb{E}[f]}{k - i} \prod_{j=i}^{k-2}(1 - \gamma_j) - c_{k,i,\gamma}\varepsilon.$$

where $c_{k,i,\gamma} \leq (k - i + 1)\left(\frac{1}{k-i}\prod_{j=i}^{k-2}(1 - \gamma_j) - \frac{1}{k-i+1}\prod_{j=i-1}^{k-2}(1 - \gamma_j)\right)\left(1 + (k - i)\gamma_{i-1}^{(-)}\right)^{-1}$

For many regimes of interest, e.g. when either γ_{k-2} or $|\gamma_{k-2}^{(-)}|$ is less than roughly $\frac{1}{k}$, $c_{\gamma,k,i} \leq O(1)$ is an *absolute constant*, with no dependence on dimension. In such cases the contrapositive of Theorem 1 states *any non-expanding function has constant density in some link*. We discuss the relation of this bound to prior literature further in Section 5. It remains an interesting open problem whether the expression is tight, or could be improved to match the expansion profile known under stronger expansion guarantees.

As applications of Theorem 1, we give the first structure theorems for simpler combinatorial constructions of high dimensional expanders. We'll start by looking at one of the most classical constructions of simplicial complexes: *clique-complexes*. Given a graph $G = (V, E)$, the k-dimensional clique-complex $K_{G,k}$ is the complex induced by taking the uniform distribution over k-cliques of G. Using Theorem 1, we give the first non-trivial characterization of non-expanding sets on (dense) clique-complexes.

▶ **Corollary 2** (Expansion in Clique-Complexes). *Fix $k \in \mathbb{N}$ and let $G = (V, E)$ be any graph with minimum degree at least $\Delta_{min} \geq \frac{2k-2}{2k-1}|V|$. Then the expansion of any (ε, i)-global boolean function f on the k-dimensional clique-complex $K_{G,k}$ is at least:*

$$\Phi(f) \geq (1 - \mathbb{E}[f])\frac{(i + 1)}{k(k - i)} - O(\varepsilon)$$

While Corollary 2 is non-trivial in any dimension, it is strongest in low dimensions. For instance, an interesting implication for 3-clique (triangle) complexes is that any function of triangles with expansion worse than 1/3 must have a dense vertex.

▶ **Corollary 3** (Low Influence Functions on Triangle Complexes). *Let $G = (V, E)$ be any graph of minimum degree at least $\Delta_{min} \geq \frac{5}{6}|V|$ and $S \subset K_{G,3}(3)$ any set with expansion*

$$\Phi(S) \leq \frac{1}{3} - \delta.$$

Then there exists a vertex $v \in G$ such that S contains at least a $\frac{\delta}{2}$-fraction of the triangles touching v.

[2] Note here we really mean the set of k-faces including some fixed i-face, sometimes called a star (rather than removing the common face as in the formal definition of a link).

Of course, clique-complexes arising from dense graphs are themselves dense objects. A major selling points of prior work was the ability to characterize low-influence functions on several families of *sparse* complexes. While our work cannot match the quantitative strength of these results, it does lead to new insight on the structure of functions on simpler constructions of bounded-degree HDX with quantitatively weaker parameters. In particular, we consider (cutoffs of) the elementary combinatorial HDX of Golowich [40], which tensors a bounded-degree expander G with the complete complex to give a weak bounded-degree two-sided HDX.

▶ **Corollary 4** (Expansion in Combinatorial HDX). *Let X be a $2k$-dimensional "product-HDX" of any graph G and $\Delta_{4k}(2k)$, and $f \in C_k$ any (ε, i)-global boolean function. Then*

$$\Phi(f) \geq \frac{(1-\mathbb{E}[f])i}{(k-1)(k-i)} - O(\varepsilon)$$

As for the case of clique-complexes, this is especially powerful in low dimensions. For concreteness, we again look at the first non-trivial setting of $k = 3$.

▶ **Corollary 5** (Low Influence Functions on Combinatorial HDX). *Let X be a 3-dimensional "product-HDX" of any graph G and $\Delta_n(6)$, and $S \subset X(3)$ be any subset of triangles with expansion at most*

$$\Phi(S) \leq \frac{1}{4} - \delta.$$

Then there exists a vertex $v \in X(1)$ such that S contains a $\Omega(\delta)$ fraction of triangles including the vertex v.

Finally, we note that as an immediate corollary of Theorem 1 we also get a "Kruskal-Katona" type theorem in all the above settings. Kruskal-Katona is a seminal and broadly used result in extremal combinatorics which given a set $S \subset X(k)$, lower bounds the number of $(k-1)$-faces that sit inside S, called its "lower shadow" and denoted ∂S.

▶ **Corollary 6** (Kruskal-Katona for Weak HDX). *Let (X, Π) be a k-dimensional simplicial complex and $S \subset X(k)$ any (ε, i)-global set. Then*

$$\mathbb{E}[\partial S] \geq \mathbb{E}[S]\left(1 + \frac{1-\mathbb{E}[S]}{k-i}\prod_{j=i}^{k-2}(1-\gamma_j) - c_{k,i,\gamma}\varepsilon.\right)$$

where $c_{k,i,\gamma} \leq (k-i+1)\left(\frac{1}{k-i}\prod_{j=i}^{k-2}(1-\gamma_j) - \frac{1}{k-i+1}\prod_{j=i-1}^{k-2}(1-\gamma_j)\right)\left(1+(k-i)\gamma_{i-1}^{(-)}\right)^{-1}$

3.2 Small-Set Expansion

Similar to the setting of the lower walk, links also provide a canonical family of non-expanding sets for the noise operator on simplicial complexes. In this setting, it is typical [56, 44, 7] to use a slightly stronger notion of global functions we'll call "(ε, i)-strongly global", which simply assumes f is sparse on every link:

$$\forall \tau \in X(i) : \mathop{\mathbb{E}}_{X_\tau}[f_\tau] \leq \varepsilon.$$

BHKL [7], GLL [44], and Hopkins [47] prove small-set expansion for strongly global functions on two-sided and one-sided partite HDX with very strong quantitative expansion.[3] While these properties are satisfied by some constructions (e.g. [55]), some variants of the seminal *Ramanujan complexes* (quotients of the affine building) of Lubotzky, Samuels, and Vishne [65] actually fail these conditions.

On the other hand, these complexes exhibit nice *local* structure. In particular, for large enough fields, their links are quantitatively strong partite HDX, and therefore satisfy a small-set expansion theorem. This raises a natural question: is any complex X whose links satisfy global small-set expansion itself a global small-set expander? We answer this question in the affirmative.

▶ **Theorem 7** (Local-to-Global SSE). *Let X be a complex such that $\forall v \in X(1)$:*

1. *Any $(\varepsilon, i-1)$-strongly global function f on X_v has expansion*

$$\Phi_{T_\rho}(f) \geq \phi(\varepsilon, k, i)$$

2. *X_v is a γ-two-sided or γ-one-sided partite HDX*

Then any (ε, i)-global boolean function f has expansion at least

$$\Phi_{T_\rho}(f) \geq \phi(\varepsilon, k, i) - 2^{O(k)}\gamma - 2^{-\Omega(k)}.$$

We note the requirement that links of X are HDX can be removed up to worsened dependence in the dimensional error term (see the full version). The inverse exponential error above is essentially negligible, since any function with respect to the noise operator has non-expansion at least inverse exponential in k.

As an immediate corollary, we prove the Ramanujan complexes are global small-set expanders.

▶ **Corollary 8** (Ramanujan Complexes are Global SSEs). *Let $f \in C_k$ be an (ε, i)-strongly global boolean function on a Ramanujan complex of [65] with $2^{-\Omega(k)}$ one-sided local-spectral expansion. The expansion of f is at least*

$$\Phi_{T_\rho}(f) \geq 1 - \rho^i - O_i(\varepsilon) - 2^{-\Omega(k)}.$$

4 Techniques

We sketch the proofs of Theorem 1 and Theorem 7, starting with the former. The formal versions and proofs of all results are deferred to the full version of the paper.

4.1 Proof Overview of Theorem 1

Our general approach to prove global sets expand is broken into three core components. The first is a new function-dependent variant of "Garland's Lemma", a classical method of breaking global functions into local components over the complex. In particular, we show the expansion of a function f with respect to the lower walks can be written as an expectation of its expansion over links of X.

[3] In particular, they require local-spectral expansion to be inverse exponential in the dimension.

▶ **Lemma 9** (Garland's Lemma for Expansion). *Let (X, Π) be a weighted, pure simplicial complex and $f \in C_k$ any function on k-faces. Then for all $i \leq k$ and $j \leq k - i$:*

$$\Phi(f) = \mathop{\mathbb{E}}_{\tau \sim \pi_j^f}[\Phi_{X_\tau}(f|_\tau)], \tag{1}$$

where $f|_\tau(\sigma) = f(\tau \cup \sigma)$ is the localization of f to the link of τ, and π_j^f is some f-dependent distribution over j-faces of X.

We remark that to the authors' knowledge, this is the first variant of Garland's lemma where the expectation is over a *function-dependent* distribution, rather than simply over the complex distribution π_j, which is necessary for our setting.

The second core component is a variant of the GK's decomposition theorem [42] for the lower walk.

▶ **Theorem 10** (GK-Decomposition). *Let (X, Π) be a k-dimensional simplicial complex. Then for any $f \in C_k$ there is an orthogonal decomposition $f = \sum\limits_{i=0}^{k} f_i$ such that:*

$$\langle N_1^k f, f \rangle \leq \mathbb{E}[f]^2 + \sum_{i=1}^{k-1} \left(1 - \frac{1}{k - i + 1} \prod_{j=i-1}^{k-2} (1 - \gamma_j) \right) \langle f, f_i \rangle.$$

Finally, the third main component is a "level-1 inequality" showing any $(\varepsilon, 1)$-global function has low projection onto the first level of the GK-decomposition.

▶ **Proposition 11** (Level-1 Inequality). *Let (X, Π) be a k-dimensional complex and f any $(\varepsilon, 1)$-global boolean function $f \in C_k$. Then*

$$\langle f, f_1 \rangle \leq \frac{k}{1 + \gamma_0^{(-)}(k-1)} \varepsilon.$$

We remark that this is a corollary of a more general level i inequality that scales with the minimum non-zero eigenvalue of the upper walks (see full version). One could instead apply this directly with the GK decomposition, but the resulting dependence on ε is worse than can be achieved by combining these with Garland's method.

Given these three components, Theorem 1 follows from a careful induction. We first prove the statement directly for $(\varepsilon, 1)$-global functions using the GK-Decomposition and Level-1 inequality. For general i, we use Garland's lemma to reduce to $i = 1$ inside links of X, writing the expansion as $\Phi(f) = \mathbb{E}_{\tau \sim \pi_{i-1}^f}[\Phi_\tau(f|_\tau)]$, and observe that the localization $f|_\tau$ of a (ε, i) global function is itself $(\varepsilon - \delta_\tau, 1)$ global for some $|\delta_\tau| \leq \varepsilon$. Applying the statement for 1-global functions inside the expectation (carefully handling dependence on $\mathbb{E}_{\tau \sim \pi_{i-1}^f}[\delta_\tau]$) then gives the claimed bound.

4.2 Proof Overview of Theorem 7

The first main component of our low influence characterization, Lemma 9, already shows that expansion of the lower walk exhibits local-to-global behavior. The proof of Theorem 7 would follow immediately if the same form of localization were to hold for the noise operator. Unfortunately, while the lemma may be extended to lower walks of arbitrary length, such a statement does not hold directly for the noise operator.

To see why, recall the noise operator is a convex combination of lower walks $T_\rho = \sum_{i=0}^{k} B_k^\rho(i) N_k^{k-i}$ where $B_k^\rho(i) = \binom{k}{i}(1-\rho)^{k-i}\rho^i$. Since expansion is linear, for any function f we can write

$$\Phi_{T_\rho}(f) = \sum_{i=0}^{k} B_k^\rho(i)\Phi_{N_k^{k-i}}(f).$$

Localizing T_ρ naively hits two main issues. First, a walk N_k^i can only be appropriately localized to j-links when $j \leq k-i$ (otherwise the walk goes below the level of the link, and cannot be captured at this level of locality). Second, the *coefficients* of T_ρ, which are binomially distributed, depend on dimension and therefore do not localize correctly. In particular, after localizing "permissible" lower walks in the sum, we'd get:

$$\Phi_{T_\rho}(f) = \sum_{i=0}^{j-1} B_k^\rho(i)\Phi_{N_k^i}(f) + \mathbb{E}_{\tau\sim\pi_j^f}\left[\sum_{i=j}^{k} B_k^\rho(i)\Phi_{N_{k-j}^{k-i}}(f|_\tau).\right]$$

To correspond to the local noise operator on the right-hand side, we'd instead need the corresponding coefficients to be $B_{k-j}^\rho(i)$, the binomial distribution on $k-j$ trials, not k trials as in the original instance. A first approach to solving these issues is to define a "shifted" noise operator over the original space that simply throws out the lefthand sum, which is negligible, and replaces the remaining coefficients with $B_{k-j}^\rho(i)$ over the original space. One can then argue that the expansion with respect to the shifted operator "localizes correctly" (allowing us to apply local SSE), and is moreover close to the expansion of the original noise operator. Unfortunately, this procedure results in error scaling with the TV-distance between Binomial distributions on k and $k-j$ trials. This costs $\text{poly}(k^{-1})$ additive error, and therefore does not give the desired global SSE theorem.

To fix this issue, we instead make the stronger assumption that links of X are good two-sided or partite one-sided high dimensional expanders, a property satisfied by our main motivating example (the Ramanujan complexes). This allows us to appeal to the theory of Fourier analysis developed in [19, 6, 7, 44], which shows one may decompose expansion with respect to the lower walks as

$$\Phi_{N_k^i}(f) \approx \frac{1}{\langle f,f\rangle}\sum_{\ell=0}^{k}\lambda(N_k^i)\langle f_\ell, f_\ell\rangle$$

for some decomposition $f = \sum_{i=0}^{k} f_i$ and approximate eigenvalues $\lambda(N_k^i)$. With this in mind, instead of shifting the noise operator, we can decompose the localized expansion in the Fourier basis as

$$\Phi_{T_\rho}(f) \lesssim \sum_{i=0}^{j-1} B_k^\rho(i) + \mathop{\mathbb{E}}_{\tau\sim\pi_j^f}\left[\frac{1}{\langle f|_\tau, f|_\tau\rangle}\sum_{i=j}^{k} B_k^\rho(i)\sum_{\ell=0}^{k-j}\lambda_\ell(N_{k-j}^{k-i}|_\tau)\langle (f|_\tau)_\ell, (f|_\tau)_\ell\rangle\right]$$

Finally for each fixed ℓ, we show its corresponding coefficient in the above sum is at most ρ^ℓ, the (approximate) eigenvalue of the local noise operator itself [7, 44]. This allows us to write the desired localized inequality

$$\Phi_{T_\rho}(f) \lesssim \sum_{i=0}^{j-1} B_k^\rho(i) + \mathop{\mathbb{E}}_{\tau\sim\pi_j^f}\left[\Phi_{T_\rho}(f|_\tau)\right]$$

which may finally be bounded by the local SSE theorem using the fact that any j-restriction $f|_\tau$ of an (ε, i)-strongly global function is $(2\varepsilon, i-j)$-strongly global. The error in this analysis scales only with the cost of cutting off small i (the first term) and the approximation error of the Fourier decomposition. For constant j, the former is at most $\exp(-k)$, while the latter is $2^{O(k)}\gamma$ when j-links are γ-HDX. Thus for complexes with sufficiently expanding links, we get an SSE theorem with negligible additive error as desired.

5 Discussion and Related Work

We close the extended abstract with a more in-depth comparison with prior work and a natural open problem such comparisons raise. At a high level, our work fits into a long line of research examining the structure of low-influence functions beyond the hypercube, starting not long after the KKL Theorem itself with the study of product spaces and the p-biased cube [13, 70, 38, 36, 46, 56] and more recently in extended settings such as the slice [60], multi-slice [35, 69], symmetric group [34], Grassmannian [62, 31], high dimensional expanders [6, 7, 44, 39, 42], and Lie groups [30]. We rely heavily on the Fourier analytic machinery initiated by Dikstein, Dinur, Filmus, and Harsha [19], and Kaufman and Oppenheim [55], and further extended by Bafna, Hopkins, Kaufman, and Lovett [6, 7], Gur Lifshitz, and Liu [44], Hopkins [47], and Gotlib and Kaufman [42].

Concretely, several bounds are known for the expansion of global sets with respect to the lower walk. In the regime of strong (inverse exponential) quantitative expansion, BHKL [6] prove any (ε, i)-global function has expansion at least

$$\Phi(f) \gtrsim (1-\mathbb{E}[f])\frac{i+1}{k} - \binom{k}{i}\varepsilon.$$

This is tight when $\varepsilon \ll \frac{1}{\binom{k}{i}}$, and mimics the structure one would expect to see on general k-fold product spaces. For strongly global functions, BHKL [7], GLL [44], and Hopkins [47] improve the ε-dependence to be *independent of dimension*:

$$\Phi(f) \gtrsim (1-\mathbb{E}[f])\frac{i+1}{k} - 2^{O(i)}\varepsilon.$$

In the regime of weak high dimensional expansion, the only known bound for global functions is due to Gotlib and Kaufman [41], who show that on any *one-sided* HDX:

$$\Phi(f) \geq (1-\mathbb{E}[f])\frac{1}{k-i}\prod_{j=i}^{k-2}(1-\gamma_j) + |X(k)|\varepsilon.$$

Unlike the prior bounds, however, this is only really meaningful when $\varepsilon \approx 0$, a non-trivial but extremely strong notion of globalness that does not lead to meaningful structural characterizations.

In some sense our bound, Theorem 1, sits between the latter two. it has the beneficial properties of the latter in that it holds for global functions and under arbitrary (two-sided) local-spectral expansion, and of the former in that the dependence on ε is typically constant. On the other hand, due to being obtained through an inductive approach, it also inherits the main term of Gotlib-Kaufman [42], scaling with $\frac{1}{k-i}$ instead of $\frac{i+1}{k}$. When i and k are large, the former dependence is substantially worse.

At a technical level, this difference in main term stems from the fact that our "advantage" for i-global functions, like in Gotlib-Kaufman [42], comes from the second eigenvalue of the lower walk inside i-links of X, which scales as $1 - \frac{1}{k-i+1}\prod\limits_{j=i-1}^{k-2}(1-\gamma_j)$. On the other hand, on very strong HDX, the advantage of i-global functions comes from the ith approximate eigenvalue of the global lower walk on X itself, which scales with $1-\frac{i}{k}$. It is natural to conjecture the "correct" dependence sits between these two bounds, as $1-\frac{i}{k}\prod_{j=i-1}^{k-2}(1-\gamma_j)$, closer to the type of bound known for the second eigenvalue from [3]. Unfortunately, both the techniques of this paper and Fourier analytic techniques from prior work (which typically incur exponential error) seem far from proving such a bound.

References

1 Rudolf Ahlswede and Peter Gács. Spreading of sets in product spaces and hypercontraction of the markov operator. *The annals of probability*, pages 925–939, 1976.

2 Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani. Approximating constraint satisfaction problems on high-dimensional expanders. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 180–201. IEEE, 2019. doi:10.1109/FOCS.2019.00021.

3 Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. In *Proceedings of the 52nd annual ACM SIGACT symposium on theory of computing*, pages 1198–1211, 2020. doi:10.1145/3357713.3384317.

4 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *SIAM Journal on Computing*, 53(6):FOCS20–1, 2021.

5 Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials ii: high-dimensional walks and an FPRAS for counting bases of a matroid. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1–12, 2019. doi:10.1145/3313276.3316385.

6 Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett. High dimensional expanders: Eigenstripping, pseudorandomness, and unique games. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1069–1128. SIAM, 2022. doi:10.1137/1.9781611977073.47.

7 Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett. Hypercontractivity on high dimensional expanders. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 185–194, 2022. doi:10.1145/3519935.3520040.

8 Mitali Bafna, Noam Lifshitz, and Dor Minzer. Constant degree direct product testers with small soundness. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 862–869. IEEE, 2024. doi:10.1109/FOCS61266.2024.00059.

9 Mitali Bafna, Dor Minzer, Nikhil Vyas, and Zhiwei Yun. Quasi-linear size pcps with small soundness from hdx. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 45–53, 2025. doi:10.1145/3717823.3718197.

10 Boaz Barak, Pravesh K. Kothari, and David Steurer. Small-Set Expansion in Shortcode Graph and the 2-to-2 Conjecture. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, volume 124 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ITCS.2019.9.

11 Daniel Beaglehole, Max Hopkins, Daniel Kane, Sihan Liu, and Shachar Lovett. Sampling equilibria: Fast no-regret learning in structured games. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3817–3855. SIAM, 2023. doi:10.1137/1.9781611977554.CH149.

12 Antonio Blanca, Pietro Caputo, Zongchen Chen, Daniel Parisi, Daniel Štefankovič, and Eric Vigoda. On mixing of markov chains: Coupling, spectral independence, and entropy factorization. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3670–3692. SIAM, 2022. doi:10.1137/1.9781611977073.145.

13 Jean Bourgain, Jeff Kahn, Gil Kalai, Yitzhak Katznelson, and Nathan Linial. The influence of variables in product spaces. *Israel Journal of Mathematics*, 77(1):55–64, 1992.

14 Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Rapid mixing for colorings via spectral independence. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1548–1557. SIAM, 2021.

15 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of glauber dynamics: Entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1537–1550, 2021. doi:10.1145/3406325.3451035.

16 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of glauber dynamics up to uniqueness via contraction. *SIAM Journal on Computing*, 52(1):196–237, 2023. doi:10.1137/20M136685X.

17 Yotam Dikstein. New high dimensional expanders from covers. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 826–838, 2023. doi:10.1145/3564246.3585183.

18 Yotam Dikstein and Irit Dinur. Agreement testing theorems on layered set systems. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1495–1524. IEEE, 2019. doi:10.1109/FOCS.2019.00088.

19 Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha. Boolean Function Analysis on High-Dimensional Expanders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*, volume 116 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.APPROX-RANDOM.2018.38.

20 Yotam Dikstein, Irit Dinur, and Alexander Lubotzky. Low acceptance agreement tests via bounded-degree symplectic hdxs. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 826–861. IEEE, 2024. doi:10.1109/FOCS61266.2024.00058.

21 Yotam Dikstein and Max Hopkins. Chernoff bounds and reverse hypercontractivity on hdx. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 870–919. IEEE, 2024. doi:10.1109/FOCS61266.2024.00060.

22 Yotam Dikstein, Max Hopkins, Russell Impagliazzo, and Toniann Pitassi. High rate efficient local list decoding from hdx. *arXiv preprint*, 2026. doi:10.48550/arXiv.2601.22535.

23 Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 357–374, 2022. doi:10.1145/3519935.3520024.

24 Irit Dinur, Yuval Filmus, Prahladh Harsha, and Madhur Tulsiani. Explicit sos lower bounds from high-dimensional expanders. In *12th Innovations in Theoretical Computer Science Conference, ITCS 2021*, pages 1–16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, 2021. doi:10.4230/LIPIcs.ITCS.2021.38.

25 Irit Dinur and Tali Kaufman. High dimensional expanders imply agreement expanders. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 974–985. IEEE, 2017. doi:10.1109/FOCS.2017.94.

26 Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 940–951, 2018. doi:10.1145/3188745.3188806.

27 Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 376–389, 2018. doi:10.1145/3188745.3188804.

28 Irit Dinur, Siqi Liu, and Rachel Yun Zhang. New Codes on High Dimensional Expanders. In *40th Computational Complexity Conference (CCC 2025)*, volume 339 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:42. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.CCC.2025.27.

29 David Ellis, Guy Kindler, and Noam Lifshitz. An analogue of bonami's lemma for functions on spaces of linear maps, and 2-2 games. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 656–660, 2023. doi:10.1145/3564246.3585116.

30 David Ellis, Guy Kindler, Noam Lifshitz, and Dor Minzer. Product mixing in compact lie groups. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1415–1422, 2024. doi:10.1145/3618260.3649626.

31 David C Ellis, Guy Kindler, and Noam Lifshitz. Forbidden intersection problems for families of linear maps. *Discrete Analysis*, 19, 2023.

32 Yumou Fei, Dor Minzer, and Shuo Wang. Multi-pass streaming lower bounds for approximating max-cut. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1537–1560, 2025. doi:10.1109/FOCS63196.2025.00081.

33 Weiming Feng, Heng Guo, Yitong Yin, and Chihao Zhang. Rapid mixing from spectral independence beyond the boolean domain. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1558–1577. SIAM, 2021. doi:10.1137/1.9781611976465.95.

34 Yuval Filmus, Guy Kindler, Noam Lifshitz, and Dor Minzer. Hypercontractivity on the symmetric group. In *Forum of Mathematics, Sigma*, volume 12, page e6. Cambridge University Press, 2024.

35 Yuval Filmus, Ryan O'Donnell, and Xinyu Wu. Log-sobolev inequality for the multislice, with applications. *Electronic Journal of Probability*, 27:1–30, 2022.

36 Ehud Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):27–35, 1998. doi:10.1007/PL00009809.

37 Ehud Friedgut and Jean Bourgain. Sharp thresholds of graph properties, and the k-sat problem. *Journal of the American mathematical Society*, 12(4):1017–1054, 1999.

38 Ehud Friedgut and Gil Kalai. Every monotone graph property has a sharp threshold. *Proceedings of the American mathematical Society*, 124(10):2993–3002, 1996.

39 Jason Gaitonde, Max Hopkins, Tali Kaufman, Shachar Lovett, and Ruizhe Zhang. Eigenstripping, Spectral Decay, and Edge-Expansion on Posets. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*, volume 245 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:24. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.APPROX/RANDOM.2022.16.

40 Louis Golowich. Improved Product-Based High-Dimensional Expanders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*, volume 207 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.APPROX/RANDOM.2021.38.

41 Roy Gotlib and Tali Kaufman. Beyond the worst case: Structured convergence of high dimensional random walks. *CoRR*, 2022. doi:10.48550/arXiv.2208.03241.

42 Roy Gotlib and Tali Kaufman. Fine grained analysis of high dimensional random walks. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2023.

43 Ben Green and Mehtaab Sawhney. New bounds for the Furstenberg-Sárközy theorem. *arXiv preprint*, 2024. arXiv:2411.17448.

44 Tom Gur, Noam Lifshitz, and Siqi Liu. Hypercontractivity on high dimensional expanders. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 176–184, 2022. doi:10.1145/3519935.3520004.

45 Tom Gur, Dor Minzer, Guy Weissenberg, and Kai Zhe Zheng. 3-query rldcs are strictly stronger than 3-query ldcs. *arXiv preprint*, 2025. doi:10.48550/arXiv.2512.12960.

46 Hamed Hatami. A structure theorem for boolean functions with small total influences. *Annals of Mathematics*, pages 509–533, 2012.

47 Max Hopkins. Hypercontractivity on hdx ii: Symmetrization and q-norms. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 96–103, 2025. `doi:10.1145/3717823.3718187`.

48 Max Hopkins and Ting-Chun Lin. Explicit lower bounds against ω (n)-rounds of sum-of-squares. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 662–673. IEEE, 2022. `doi:10.1109/FOCS54457.2022.00069`.

49 Jun-Ting Hsieh, Alexander Lubotzky, Sidhanth Mohanty, Assaf Reiner, and Rachel Yun Zhang. Explicit lossless vertex expanders. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 894–911, 2025. `doi:10.1109/FOCS63196.2025.00046`.

50 Vishesh Jain, Huy Tuan Pham, and Thuy-Duong Vuong. Spectral independence, coupling, and the spectral gap of the glauber dynamics. *Information Processing Letters*, 177:106268, 2022. `doi:10.1016/j.ipl.2022.106268`.

51 J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 68–80, Los Alamitos, CA, USA, October 1988. IEEE Computer Society. `doi:10.1109/SFCS.1988.21923`.

52 Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Isoperimetric inequalities for ramanujan complexes and topological expanders. *Geometric and Functional Analysis*, 26(1):250–287, 2016.

53 Tali Kaufman and David Mass. Local-To-Global Agreement Expansion via the Variance Method. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 74:1–74:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.ITCS.2020.74`.

54 Tali Kaufman and Izhar Oppenheim. Construction of new local spectral high dimensional expanders. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 773–786, 2018. `doi:10.1145/3188745.3188782`.

55 Tali Kaufman and Izhar Oppenheim. High order random walks: Beyond spectral gap. *Combinatorica*, pages 1–37, 2020.

56 Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer. Hypercontractivity for global functions and sharp thresholds. *Journal of the American Mathematical Society*, 37(1):245–279, 2024.

57 Peter Keevash, Noam Lifshitz, and Dor Minzer. On the largest product-free subsets of the alternating groups. *Inventiones mathematicae*, 237(3):1329–1375, 2024.

58 Nathan Keller, Noam Lifshitz, and Omri Marcus. Sharp hypercontractivity for global functions. *Journal of the European Mathematical Society*, 2026. Published online first. `doi:10.4171/JEMS/1762`.

59 Nathan Keller, Noam Lifshitz, Dor Minzer, and Ohad Sheinfeld. On t-intersecting families of permutations. *Advances in Mathematics*, 445:109650, 2024.

60 Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra. Small-set expansion in the johnson graph. *Theory of Computing*, 21(1):1–43, 2025. `doi:10.4086/TOC.2025.V021A002`.

61 Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and grassmann graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 576–589, 2017. `doi:10.1145/3055399.3055432`.

62 Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 592–601. IEEE, 2018. `doi:10.1109/FOCS.2018.00062`.

63 Noam Lifshitz and Dor Minzer. Noise sensitivity on the p-biased hypercube. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1205–1226. IEEE, 2019. `doi:10.1109/FOCS.2019.00075`.

64 Kuikui Liu. From Coupling to Spectral Independence and Blackbox Comparison with the Down-Up Walk. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*, volume 207 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 32:1–32:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. `doi:10.4230/LIPIcs.APPROX/RANDOM.2021.32`.

65 Alexander Lubotzky, Beth Samuels, and Uzi Vishne. Explicit constructions of ramanujan complexes of type Ad. *European Journal of Combinatorics*, 26(6):965–993, 2005. doi:10.1016/J.EJC.2004.06.007.

66 Dor Minzer and Kai Zhe Zheng. Near optimal alphabet-soundness tradeoff pcps. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 15–23, 2024. doi:10.1145/3618260.3649606.

67 Izhar Oppenheim. Local spectral expansion approach to high dimensional expanders part i: Descent of spectral gaps. *Discrete & Computational Geometry*, 59(2):293–330, 2018. doi:10.1007/S00454-017-9948-X.

68 Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical ldpc codes. In *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing*, pages 375–388, 2022. doi:10.1145/3519935.3520017.

69 Justin Salez. A sharp log-sobolev inequality for the multislice. *Annales Henri Lebesgue*, 4:1143–1161, 2021.

70 Michel Talagrand. On Russo's approximate zero-one law. *The Annals of Probability*, pages 1576–1587, 1994.