Electronic Colloquium on Computational Complexity, Report No. 188 (2026)

An elementary proof of the Komlós conjecture

Sankeerth Rao Karingula
Agentin AI
sankeerthrao@agentin.ai

Shachar Lovett*
University of California, San Diego
slovett@ucsd.edu

September 17, 2026

Abstract

We give an elementary proof of the Komlós conjecture by simplifying the recent proof of Guo, Fang, and Lu. We show that any vectors $v_1, \ldots, v_n \in \mathbb{R}^d$ with $\|v_i\|_2 \leq 1$ admit signs $\varepsilon_i \in \{-1, 1\}$ such that $\|\sum_{i=1}^n \varepsilon_i v_i\|_\infty \leq 36$. The proof uses only elementary combinatorial and probabilistic arguments and basic calculus.

1 Introduction

The Komlós conjecture is a longstanding problem in discrepancy theory, recorded by Spencer [13]. It generalizes the $O(\sqrt{n})$ discrepancy bound of Spencer [12] and contains the conjecture of Beck and Fiala [8] as a special case.

We write $\|\cdot\|_2$ and $\|\cdot\|_\infty$ for the Euclidean and maximum norms, respectively.

Conjecture 1.1 (Komlós conjecture)**.** *There is a universal constant $C > 0$ such that the following holds for all positive integers n, d. Let $v_1, \ldots, v_n \in \mathbb{R}^d$ satisfy $\|v_i\|_2 \leq 1$. Then there exist signs $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, 1\}$ such that*

$$\left\| \sum_{i=1}^n \varepsilon_i v_i \right\|_\infty \leq C.$$

Progress toward the conjecture. Banaszczyk [2] established the bound $O(\sqrt{\log n})$ through his vector-balancing theorem. Subsequent algorithmic progress includes the Gram–Schmidt walk of Bansal, Dadush, Garg, and Lovett [4], which gives a constructive version of Banaszczyk's vector-balancing theorem. Bansal and Jiang [5, 6] obtained the first asymptotic improvement over Banaszczyk's bound, proving $O((\log n)^{1/4}(\log \log n)^{7/4})$. Ercan [9] subsequently removed the iterated-logarithmic factor, obtaining $O((\log n)^{1/4})$.

Recent breakthrough. In September 2026, Guo, Fang, and Lu [10] gave a proof of the Komlós conjecture with $C = 3\sqrt{2\pi}$. The authors credit the Odin Automatic AI Research Agent with discovering the proof. Their argument uses Banaszczyk's convex-body transform and a variational analysis of the directional total variation of probability densities. The present work grew out of an effort to understand their argument and extract an elementary proof.

*Supported by Simons Investigator Award #929894 and NSF award CCF-2425349.

ISSN 1433-8092

Our results. We give an elementary proof of the Komlós conjecture, combining a combinatorial balancing argument with a continuous construction followed by discretization. The proof does not require advanced convex geometry or analysis. Our proof yields the constant 36, larger than $3\sqrt{2\pi}$ in [10]. We have prioritized clarity of exposition and made no attempt to optimize the constant. We note that our proof gives a finite construction for rational inputs, but does not establish a polynomial-time algorithm.

Theorem 1.2. *Let $v_1, \dots, v_n \in \mathbb{R}^d$ satisfy $\|v_i\|_2 \le 1$. Then there exist signs $\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}$ such that*

$$\left\| \sum_{i=1}^{n} \varepsilon_i v_i \right\|_\infty \le 36.$$

Notation. For a nonnegative function P on a domain X, write $\mathrm{supp}(P) := \{x \in X : P(x) > 0\}$ for its support. In the discrete parts of the proof, all probability distributions have finite support and total mass $\sum_x P(x) = 1$. We write $\Pr_P[A] := \sum_{x \in A} P(x)$ for the probability of an event A, and $\mu(P) := \sum_x P(x)x$ for the mean when $X = \mathbb{R}^d$. For $S \subseteq \mathbb{R}^d$, $\mathrm{conv}(S)$ denotes its convex hull. In particular, $\mu(P) \in \mathrm{conv}(\mathrm{supp}(P))$.

Proof approach. Our proof is based on distributions that are nearly invariant under prescribed translations. The central quantity is their *shift distance*, which we now define.

For finitely supported probability distributions P, Q on the same domain, their statistical distance (or total variation distance) is $d_{\mathrm{TV}}(P, Q) := \frac{1}{2} \sum_x |P(x) - Q(x)|$.

Definition 1.3 (Shift distance)**.** *Let P be a finitely supported probability distribution on $\mathbb{R}^d$, and let $u \in \mathbb{R}^d$. Write $P + u$ for the distribution of $X + u$ when $X \sim P$; thus $(P + u)(x) = P(x - u)$. The* shift distance *of P at u is*

$$\Delta(P, u) := d_{\mathrm{TV}}(P, P + u) = \frac{1}{2} \sum_{x \in \mathbb{R}^d} |P(x) - P(x - u)|.$$

Thus, small shift distance means that P changes little under the specified translation. Equivalently, since

$$\sum_{x \in \mathbb{R}^d} \min\{P(x), P(x - u)\} = 1 - \Delta(P, u),$$

$\Delta(P, u) \le \delta$ means that P and its translate have at least $1 - \delta$ of their probability mass in common. This overlap interpretation is what makes shift distance useful in the splitting argument below.

The proof has two main parts. First, we show that the mean of a distribution with small shift distances at $6v_1, \dots, 6v_n$ can be shifted by a signed sum of the v_i while remaining in the convex hull of its support. We allow an arbitrary mean so that this statement can be applied inductively after adding a coordinate. Second, we construct a mean-zero distribution with the required shift bounds in a cube by rounding a continuous product density to a rational grid.

Lemma 1.4 (From near invariance to signed sums)**.** *Let $v_1, \dots, v_n \in \mathbb{R}^d$. Suppose P is a finitely supported probability distribution on $\mathbb{R}^d$ with*

$$\Delta(P, 6v_i) \le \frac{1}{3} \qquad (i = 1, \dots, n).$$

Then there are signs $\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}$ such that

$$\mu(P) + \sum_{i=1}^{n} \varepsilon_i v_i \in \mathrm{conv}(\mathrm{supp}(P)).$$

Lemma 1.5 (A near-invariant distribution in the cube)**.** *Let $v_1, \dots, v_n \in \mathbb{Q}^d$ satisfy $\|v_i\|_2 \le 1$. There is a finitely supported distribution P on $[-6, 6]^d$ with mean zero and*

$$\Delta(P, v_i) \le \frac{1}{3} \qquad (i = 1, \dots, n).$$

Proof of Theorem 1.2. First suppose the vectors are rational. Take P from Lemma 1.5 and apply Lemma 1.4 to the vectors $v_i/6$. Since P has mean zero, this gives signs such that

$$\frac{1}{6}\sum_{i=1}^{n} \varepsilon_i v_i \in \operatorname{conv}(\operatorname{supp}(P)) \subseteq [-6,6]^d,$$

and hence $\|\sum_i \varepsilon_i v_i\|_\infty \leq 36$. For real vectors, approximate each v_i by rational vectors in the set $\{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$, for instance by truncating its coordinates toward zero. Since there are only finitely many sign vectors, some fixed choice of signs works for an infinite subsequence of these approximations. Passing to the limit proves the same bound. □

Our theorem controls the final signed sum. The *strong Komlós conjecture* (also called the prefix Komlós conjecture) asks whether one can choose signs that simultaneously keep every prefix sum bounded by a universal constant. We do not know whether our approach extends to the strong Komlós conjecture.

Conjecture 1.6 (Strong Komlós conjecture)**.** *There is a universal constant $C > 0$ such that every sequence $v_1, \ldots, v_n \in \mathbb{R}^d$ with $\|v_i\|_2 \leq 1$ admits signs $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, 1\}$ satisfying*

$$\max_{1\leq k\leq n} \left\| \sum_{i=1}^{k} \varepsilon_i v_i \right\|_\infty \leq C.$$

The order is fixed, and the signs may depend on the entire sequence. Prefix balancing goes back to Spencer [11]; the constant bound was formulated explicitly by Bansal, Jiang, Meka, Singla, and Sinha [7, Open Problem 6.1 in the full version]. Guo, Fang, and Lu [10, Section 1.2] also note that their proof does not control all prefixes. The best known bound uniform in d remains $O(\sqrt{\log n})$, due to Banaszczyk [3].[1] Aden-Ali [1] gives a randomized online algorithm that, for every input sequence fixed independently of the algorithm's random choices, chooses each sign as its vector arrives and achieves $O(\sqrt{\log n})$ prefix discrepancy with high probability.

AI methodology. We used ChatGPT-6 Astra to iteratively refine the proof and the writing, with an emphasis on simplifying the argument and improving the clarity of the exposition.

Acknowledgements. We thank Nikhil Bansal for helpful feedback and suggestions on an earlier draft.

Paper organization. We give an overview of both lemmas in Section 2, then prove Lemma 1.4 in Section 3 and Lemma 1.5 in Section 4.

2 Proof overview

2.1 From near invariance to signed sums

We prove Lemma 1.4 by induction on the number of vectors, simultaneously in all dimensions. At each step, we split once in the direction of the last vector, apply induction to the remaining vectors, and pull the resulting convex combination back to the original support.

[1] To remove the dimension dependence, let $r_1, \ldots, r_d \in \mathbb{R}^n$ be the rows of the matrix with columns $v_1, \ldots, v_n$. A row with $\|r_k\|_1 \leq 1$ has discrepancy at most 1 for every signing and every prefix. Each remaining row satisfies $\|r_k\|_2^2 \geq \|r_k\|_1^2/n > 1/n$. Since $\sum_{k=1}^{d} \|r_k\|_2^2 = \sum_{i=1}^{n} \|v_i\|_2^2 \leq n$, fewer than n^2 rows remain. Restricting to these rows does not increase any column's Euclidean norm. Applying the $O(\sqrt{\log d + \log n})$ prefix bound of Banaszczyk [3] in this reduced dimension gives $O(\sqrt{\log n})$ for $n \geq 2$; see also [7, Theorem 1.1]. If no rows remain, or if $n = 1$, the discrepancy is at most 1.

One split. Let $v = v_n$. We split P to form a distribution Q on $\mathbb{R}^d \times \{0,1\}$. A new state (x,b) has possible parents $x+3v$ and $x-3v$ in P; a parent is available if it has positive mass in P. At each position x, the split assigns half the larger parent mass to $(x,0)$ and half the smaller parent mass to $(x,1)$. Every state of positive mass therefore has an available parent, and a state with $b=1$ has both.

We regard the bit b as an extra real coordinate. The first d coordinates retain their mean $\mu(P)$. The last coordinate has mean $\beta := \Pr_Q[b=1]$, which is half the overlap of the two translates of P. Since these translates differ by $6v$, the assumption $\Delta(P,6v) \le 1/3$ gives $\beta \ge 1/3$. Thus $\mu(Q) = (\mu(P), \beta)$.

Induction. Extend each remaining vector to $(v_i, 0) \in \mathbb{R}^{d+1}$. Splitting does not increase shift distance in these directions, so induction applies to Q. It gives signs for the first $n-1$ vectors and a point

$$(\mu(P)+s, \beta) \in \operatorname{conv}(\operatorname{supp}(Q)), \qquad s := \sum_{i=1}^{n-1} \varepsilon_i v_i.$$

Represent this point as the mean of a probability distribution R supported on $\operatorname{supp}(Q)$. Its last coordinate forces $\Pr_R[b=1] = \beta \ge 1/3$. This is why we keep the extra coordinate when applying induction: it ensures that the chosen convex combination still has enough mass with two available parents.

Pullback. Send each entry of R with $b=0$ to any available parent in P. These entries have total mass at most two-thirds, and each move is by $\pm 3v$. Their contribution to the change in the first d coordinates of the mean is therefore av for some $|a| \le 2$. The entries with $b=1$ have both parents available and total mass at least one-third. By dividing this mass between its parents, we can contribute any scalar multiple of v with coefficient between -1 and 1. Choose $\varepsilon_n \in \{-1,1\}$ with $|\varepsilon_n - a| \le 1$ and use this mass to contribute $(\varepsilon_n - a)v$. The resulting distribution is supported on $\operatorname{supp}(P)$ and has mean $\mu(P) + s + \varepsilon_n v$, completing the induction step.

Connection to the original proof. Guo, Fang, and Lu [10, Section 4] introduce an extra coordinate and rearrange a density along it, placing larger values closer to zero. This preserves mass and does not increase total variation distance between translates in the original coordinates. Our splitting operator uses a bit as the extra coordinate and sorts two translated probabilities. Preserving the mean of this coordinate under induction supplies the mass needed for the pullback. This simplifies the argument, at the cost of a larger constant in the final bound.

2.2 A near-invariant distribution in the cube

For Lemma 1.5, we first construct a continuous probability density F on $[-6,6]^d$. We take F to be a product of one-dimensional densities, each proportional to $\max\{6-|t|, 0\}^2$. To estimate the effect of translation, we work with $f = \sqrt{F}$, a product of tent functions. When we square the derivative of f in direction v and integrate, the mixed terms vanish by symmetry, leaving a multiple of $\sum_k v_k^2 = \|v\|_2^2$. The fundamental theorem of calculus and Cauchy–Schwarz then give total variation distance at most $\|v\|_2/\sqrt{12}$ for a translation by v.

We then discretize the construction to make it finitely supported. We choose a positive integer N such that the grid $N^{-1}\mathbb{Z}^d$ contains all the given rational vectors, and round each coordinate to the nearest grid point. Since ± 6 are grid points, rounding keeps the support in the cube. It also preserves symmetry, giving a finitely supported distribution with mean zero. Rounding commutes with translations by grid vectors and cannot increase total variation distance. The same bound therefore holds for the rounded distribution. Since $\|v_i\|_2 \le 1$, this bound is less than $1/3$ for each given vector.

3 From near invariance to signed sums

The proof of Lemma 1.4 uses one split, an application of induction, and one pullback. The split adds a coordinate taking values in $\{0,1\}$, so we view its output as a distribution on $\mathbb{R}^{d+1}$. For $u \in \mathbb{R}^d$, write $\bar{u} := (u, 0)$ for the corresponding vector in $\mathbb{R}^{d+1}$.

Definition 3.1 (Splitting operator)**.** *Let $v \in \mathbb{R}^d$, and let P be a finitely supported nonnegative function on $\mathbb{R}^d$. Define $T_v P$ on $\mathbb{R}^{d+1}$ by*

$$(T_v P)(x, 0) := \frac{1}{2} \max\{P(x+v), P(x-v)\},$$
$$(T_v P)(x, 1) := \frac{1}{2} \min\{P(x+v), P(x-v)\},$$

and set it to zero whenever the last coordinate is not 0 *or* 1*.*

We call $x+v$ and $x-v$ the two possible parents of (x,b). A parent is available if it belongs to $\operatorname{supp}(P)$. Every state in $\operatorname{supp}(T_v P)$ has at least one available parent, and a state with $b=1$ has both. We will use this freedom in the pullback.

Mass and mean. Adding the maximum and minimum gives

$$\sum_{b=0}^{1} (T_v P)(x,b) = \frac{1}{2}\big(P(x+v) + P(x-v)\big). \tag{1}$$

Summing over x shows that T_v preserves total mass, even for unnormalized nonnegative inputs. In particular, it maps probability distributions to probability distributions.

Now suppose P is a probability distribution. The two opposite shifts cancel in the mean of the first d coordinates:

$$\sum_{x,b} (T_v P)(x,b)x = \frac{1}{2}\sum_x (x-v)P(x) + \frac{1}{2}\sum_x (x+v)P(x) = \mu(P).$$

The mean of the last coordinate is its probability of being 1. By the overlap interpretation of shift distance,

$$\begin{aligned}
\Pr_{(x,b)\sim T_v P}[b=1] &= \frac{1}{2}\sum_x \min\{P(x+v), P(x-v)\} \\
&= \frac{1}{2}\sum_x \min\{P(x), P(x-2v)\} \\
&= \frac{1-\Delta(P,2v)}{2}.
\end{aligned} \tag{2}$$

Thus, writing $\beta := \Pr_{T_v P}[b=1]$, we have

$$\mu(T_v P) = (\mu(P), \beta).$$

Claim 3.2 (Shift distance does not increase)**.** *Let P be a finitely supported probability distribution on $\mathbb{R}^d$, and let $u, v \in \mathbb{R}^d$. Then*

$$\Delta(T_v P, \bar{u}) \le \Delta(P, u). \tag{3}$$

Proof. Both maximum and minimum are nondecreasing in each argument, so T_v preserves pointwise inequalities between nonnegative functions. The definition also gives $T_v(P+u) = T_v P + \bar{u}$.

Consider the mass common to P and $P+u$, given by $H(x) := \min\{P(x), P(x-u)\}$. Since $H \le P$ and $H \le P+u$ pointwise, applying T_v to these inequalities gives

$$T_v H \le \min\{T_v P, T_v(P+u)\} = \min\{T_v P, T_v P + \bar{u}\}.$$

Since T_v preserves the total mass of H,

$$\begin{aligned}1 - \Delta(T_v P, \bar{u}) &= \sum_{x,b} \min\{(T_v P)(x,b), (T_v P + \bar{u})(x,b)\} \\ &\geq \sum_{x,b} (T_v H)(x,b) \\ &= \sum_{x} H(x) = 1 - \Delta(P, u).\end{aligned}$$

This proves (3). $\square$

Proof of Lemma 1.4. We induct on n, proving the statement simultaneously for all dimensions d. For $n = 0$, the conclusion is $\mu(P) \in \operatorname{conv}(\operatorname{supp}(P))$, which holds because $\mu(P)$ is a convex combination of the support points.

Suppose $n \geq 1$, let $v = v_n$, and set $Q := T_{3v} P$. By (2),

$$\beta := \Pr_{(x,b)\sim Q}[b = 1] = \frac{1 - \Delta(P, 6v)}{2} \geq \frac{1}{3}, \qquad \mu(Q) = (\mu(P), \beta).$$

For $i < n$, Claim 3.2 gives $\Delta(Q, 6\bar{v}_i) \leq \Delta(P, 6v_i) \leq 1/3$. Apply the induction hypothesis to Q and $\bar{v}_1, \dots, \bar{v}_{n-1} \in \mathbb{R}^{d+1}$. It yields signs $\varepsilon_1, \dots, \varepsilon_{n-1}$ such that

$$(\mu(P) + s, \beta) \in \operatorname{conv}(\operatorname{supp}(Q)), \qquad s := \sum_{i=1}^{n-1} \varepsilon_i v_i.$$

Choose a probability distribution R supported on $\operatorname{supp}(Q)$ whose mean is this point. Since the last coordinate takes only values 0 and 1, its mean determines its probability of being 1:

$$\Pr_{(x,b)\sim R}[b = 1] = \beta \geq \frac{1}{3}.$$

We now pull R back to $\operatorname{supp}(P)$ by moving each state (x, b) to its available parents $x + 3v$ and $x - 3v$. For the states with $b = 0$, choose any available parent and send the entire mass there. Their total mass is $1 - \beta$, so their contribution to the change in the first d coordinates of the mean is av for some scalar a satisfying

$$|a| \leq 3(1 - \beta) \leq 2.$$

Every state with $b = 1$ has both parents available. Send a common fraction $t \in [0, 1]$ of each such state's mass to $x + 3v$ and the remaining fraction $1 - t$ to $x - 3v$. Since these states have total mass β, their contribution to the change in the mean is

$$3\beta t v - 3\beta(1 - t)v = 3\beta(2t - 1)v.$$

As t varies from 0 to 1, the coefficient ranges over $[-3\beta, 3\beta]$. Choose $\varepsilon_n \in \{-1, 1\}$ to have the sign of a, taking $\varepsilon_n = 1$ when $a = 0$. Since $a \in [-2, 2]$, we have

$$|\varepsilon_n - a| \leq 1 \leq 3\beta.$$

Set

$$t := \frac{1}{2} + \frac{\varepsilon_n - a}{6\beta}.$$

The preceding inequality ensures that $t \in [0, 1]$, and $3\beta(2t - 1) = \varepsilon_n - a$. Thus the one-bit states contribute exactly $(\varepsilon_n - a)v$. Together with the contribution av from the zero-bit states, the total change in the first d coordinates of the mean is exactly $\varepsilon_n v$.

Let S be the resulting distribution on $\mathbb{R}^d$, adding contributions that reach the same parent. Every state transfers its entire mass to available parents, so S is a probability distribution supported on $\operatorname{supp}(P)$. Its mean is

$$\mu(S) = \mu(P) + s + \varepsilon_n v = \mu(P) + \sum_{i=1}^{n} \varepsilon_i v_i \in \operatorname{conv}(\operatorname{supp}(P)),$$

as required. □

4 A near-invariant distribution in the cube

We first prove a continuous version of Lemma 1.5, using a product density whose square root has small directional derivatives. We then deduce Lemma 1.5 by rounding this density to a grid.

A probability density is a nonnegative integrable function $F : \mathbb{R}^d \to \mathbb{R}$ with $\int_{\mathbb{R}^d} F(x)\,dx = 1$. We say that F is supported on S if it vanishes outside S. For two densities, statistical distance is defined by replacing the sum with an integral: $d_{\mathrm{TV}}(F, H) := \frac{1}{2}\int_{\mathbb{R}^d} |F(x) - H(x)|\,dx$. We write $\|h\|_{L^2} := (\int_{\mathbb{R}^d} |h(x)|^2\,dx)^{1/2}$ for the L^2 norm, ∇h for the gradient, and $v \cdot \nabla h$ for the directional derivative along v.

Lemma 4.1 (A continuous near-invariant distribution in the cube)**.** *For every positive integer d, there is a continuous probability density F supported on $[-6, 6]^d$, symmetric about the origin, such that for every $v \in \mathbb{R}^d$,*

$$\frac{1}{2}\int_{\mathbb{R}^d} |F(x+v) - F(x)|\,dx \le \frac{\|v\|_2}{\sqrt{12}}.$$

Proof. Define

$$b(t) := \frac{1}{12}\max\{6 - |t|, 0\}, \qquad f(x) := \prod_{k=1}^{d} b(x_k), \qquad F(x) := f(x)^2.$$

The elementary identities

$$\int_{\mathbb{R}} b^2 = 1, \qquad \int_{\mathbb{R}} (b')^2 = \frac{1}{12}, \qquad \int_{\mathbb{R}} bb' = 0$$

show that F is a continuous probability density supported on $[-6, 6]^d$. Since b is even, $F(-x) = F(x)$, so the density is symmetric about the origin and has mean zero. These identities also give, for every $v \in \mathbb{R}^d$,

$$\int_{\mathbb{R}^d} (v \cdot \nabla f(x))^2\,dx = \frac{\|v\|_2^2}{12}. \tag{4}$$

Indeed, the mixed terms vanish because they contain the factor $\int bb' = 0$, while the kth diagonal term is $v_k^2/12$.

Because b is continuous and linear on each of finitely many intervals, the fundamental theorem of calculus along line segments gives, for almost every x,

$$f(x+v) - f(x) = \int_0^1 v \cdot \nabla f(x + tv)\,dt.$$

Values of the derivatives at the breakpoints do not affect any of the integrals. Cauchy–Schwarz therefore gives

$$\begin{aligned}\int_{\mathbb{R}^d} |f(x+v) - f(x)|^2\,dx &\le \int_0^1 \int_{\mathbb{R}^d} |v \cdot \nabla f(x+tv)|^2\,dx\,dt \\ &= \frac{\|v\|_2^2}{12}.\end{aligned}$$

Factoring a difference of squares and applying Cauchy–Schwarz once more, we obtain

$$\begin{aligned}\frac{1}{2}\int_{\mathbb{R}^d} |F(x+v)-F(x)|\,dx &\le \frac{1}{2}\|f(\cdot+v)-f\|_{L^2}\|f(\cdot+v)+f\|_{L^2} \\ &\le \|f(\cdot+v)-f\|_{L^2} \le \frac{\|v\|_2}{\sqrt{12}},\end{aligned} \tag{5}$$

where we used $\|f\|_{L^2} = \|f(\cdot+v)\|_{L^2} = 1$. ☐

We now round the continuous density to obtain the finitely supported distribution required by Lemma 1.5.

Proof of Lemma 1.5. Let F be the symmetric density from Lemma 4.1. Choose a positive integer N such that $Nv_i \in \mathbb{Z}^d$ for every i, and let $G = N^{-1}\mathbb{Z}^d = \{k/N : k \in \mathbb{Z}^d\}$ be the resulting grid. Round each coordinate of a sample from F to the nearest multiple of $1/N$. The resulting distribution is

$$P(x) := \int_C F(x+z)\,dz \quad (x \in G), \qquad C := \left[-\frac{1}{2N}, \frac{1}{2N}\right)^d,$$

and $P(x) = 0$ outside G. Because -6 and 6 are grid points, rounding keeps the support in $[-6,6]^d$, so P has finite support. Symmetry gives $P(-x) = P(x)$ and hence $\mu(P) = 0$; cell boundaries have zero probability.

Rounding commutes with grid translations and cannot increase total variation. Explicitly, for $v \in G$,

$$\begin{aligned}\Delta(P,v) &= \frac{1}{2}\sum_{x\in G}\left|\int_C \big(F(x+z-v)-F(x+z)\big)\,dz\right| \\ &\le \frac{1}{2}\int_{\mathbb{R}^d} |F(z-v)-F(z)|\,dz \le \frac{\|v\|_2}{\sqrt{12}}.\end{aligned}$$

The first inequality uses that the cells $x+C$, $x \in G$, partition $\mathbb{R}^d$; the last uses Lemma 4.1 with $-v$. Applying this to each v_i yields $\Delta(P,v_i) \le 1/\sqrt{12} < 1/3$, as required. ☐

References

[1] I. Aden-Ali. Optimal online discrepancy minimization in linear time. *arXiv preprint arXiv:2607.04388*, 2026.

[2] W. Banaszczyk. Balancing vectors and Gaussian measures of n-dimensional convex bodies. *Random Structures & Algorithms*, 12(4):351–360, 1998.

[3] W. Banaszczyk. On series of signed vectors and their rearrangements. *Random Structures & Algorithms*, 40(3):301–316, 2012.

[4] N. Bansal, D. Dadush, S. Garg, and S. Lovett. The Gram–Schmidt walk: A cure for the Banaszczyk blues. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 587–597. ACM, 2018.

[5] N. Bansal and H. Jiang. Decoupling via affine spectral-independence: Beck–Fiala and Komlós bounds beyond Banaszczyk. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, STOC 2026, pages 432–442. ACM, 2026. Preprint first posted in August 2025.

[6] N. Bansal and H. Jiang. An exposition of the $\widetilde{O}(\log^{1/4} n)$ bound for the Komlós problem, Aug. 2026. arXiv preprint arXiv:2608.28452.

[7] N. Bansal, H. Jiang, R. Meka, S. Singla, and M. Sinha. Prefix discrepancy, smoothed analysis, and combinatorial vector balancing. In M. Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. Full version: arXiv:2111.07049.

[8] J. Beck and T. Fiala. "Integer-making" theorems. *Discrete Applied Mathematics*, 3(1):1–8, 1981.

[9] E. Ercan. A $(\log n)^{1/4}$ bound for the Komlós problem, Sept. 2026. arXiv preprint arXiv:2609.08885, posted September 8, 2026.

[10] S. Guo, E. X. Fang, and J. Lu. Vector balancing via directional total variation. *arXiv preprint arXiv:2609.11189*, 2026.

[11] J. Spencer. Balancing games. *Journal of Combinatorial Theory, Series B*, 23(1):68–74, 1977.

[12] J. Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289(2):679–706, 1985.

[13] J. Spencer. *Ten Lectures on the Probabilistic Method*. Society for Industrial and Applied Mathematics, Philadelphia, 1987.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il