A modification of Alweiss's example: nilrecurrence without fixed-density largeness at any density below one

Produced by ChatGPT

15 September 2026

Abstract

We give a diagonal modification of Alweiss's construction of a set of nilrecurrence which is not a set of double topological recurrence. The modified set is, in addition, not δ-density large for any $0 < \delta < 1$, where density means natural upper density exactly δ. The same set of differences works for all these densities. The two changes are a fixed irrational distinguished coordinate and stage-dependent tail tolerances tending to zero. An elementary sparse-vector lemma, together with the uniform finite-orbit argument in Green's exposition, preserves nilrecurrence at every stage. Alweiss's complex second-difference obstruction remains uniform. For the density obstruction, sparse coarse barriers bound the tail on a long progression; its vanishing curvature then makes it nearly constant on a prescribed initial segment. The irrational distinguished coordinate forces that segment into an arbitrarily thin forbidden band. Phase averaging and a periodic deletion handle density and finitely many initial exceptions. We include all construction and approximation details; the quoted nilmanifold inputs are stated explicitly.

Contents

1 The result and the two modifications **2**

2 Sparse templates close to large subspaces **3**

3 The uniform finite nilmanifold input **4**

4 Constructing one set which retains nilrecurrence **5**

5 One fixed obstruction to double topological recurrence **8**

6 Vanishing tail curvature and narrow forbidden bands **9**

7 All densities: phase averaging, early stages, and exact thinning **12**

8 Why the diagonalization is legitimate **13**

1 The result and the two modifications

We write $\mathbb{N} = \{1, 2, \ldots\}$, $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, and

$$\|t\|_{\mathbb{T}} = \operatorname{dist}(t, \mathbb{Z}), \qquad \overline{d}(A) = \limsup_{N\to\infty} \frac{|A \cap [1, N]|}{N}.$$

All common differences of arithmetic progressions are positive.

Definition 1.1. For $0 < \delta < 1$, a set $D \subseteq \mathbb{N}$ is *δ-density large* if every $A \subseteq \mathbb{N}$ with $\overline{d}(A) = \delta$ contains arbitrarily long arithmetic progressions whose common difference belongs to D.

Theorem 1.2. *There is a single infinite set $D \subseteq \mathbb{N}$ with the following properties.*

(i) ***Multiple pointwise nilrecurrence.*** *For every finite-step nilsystem (X, T), every $x \in X$, every $\ell \geq 1$, every neighborhood U of x, and every $M \geq 1$, there is $d \in D$ with $d > M$ such that*

$$x, T^d x, T^{2d} x, \ldots, T^{\ell d} x \in U.$$

(ii) ***Failure of double topological recurrence.*** *There is one finite coloring of $\mathbb{Z}$ with no monochromatic progression $n, n + d, n + 2d$ for any $d \in D$. In particular there is a compact minimal system (Y, S) and a nonempty clopen set $V \subseteq Y$ such that*

$$V \cap S^{-d} V \cap S^{-2d} V = \varnothing \qquad (d \in D).$$

(iii) ***Failure at every fixed density below one.*** *For every $0 < \delta < 1$ there are $A_\delta \subseteq \mathbb{N}$ and an integer $K_\delta \geq 3$ such that $\overline{d}(A_\delta) = \delta$ and A_δ contains no K_δ-term arithmetic progression with common difference in D.*

Alweiss's construction has one coordinate near a fixed negative number $-a$, while the remaining coordinates are individually tiny and have total absolute mass near a; its complex exponential coloring obstructs three-term progressions [1, Definition 2.1 and Section 2.2]. We keep that architecture but make two changes:

$$a = \frac{\sqrt{2}}{1000} \text{ is irrational, and the tail scale tends to zero with the stage.}$$

The irrationality permits arbitrarily fine sampling of a fixed cosine. The shrinking tail scale supplies the increasingly accurate local control required when δ approaches one. A fixed positive value of a still gives a uniform obstruction to three-term monochromatic progressions.

Remark 1.3 (Quantifiers and scope)**.** The set D, the number a, and all frequencies used to define it are fixed before δ is specified. Only the witnessing set A_δ and the forbidden length K_δ depend on δ. This is a modification of the original example, not a proof of the corresponding assertion for Alweiss's unchanged displayed parameters. We do not claim that the density witnesses in this proof have natural density or uniform density; their natural upper density is exactly the required value.

2 Sparse templates close to large subspaces

Let

$$E = \ell^1(\mathbb{N};\mathbb{T}), \qquad \|x\|_E = \sum_{r\ge 1} \|x_r\|_{\mathbb{T}}.$$

For finite-dimensional vectors we use the same notation for the sum of the coordinate distances on a torus. The standard coordinate vectors in $\mathbb{R}^m$ are denoted by $e_1, \dots, e_m$.

A template has the form

$$v = -ae_i + b\sum_{r\in J} e_r, \qquad |J| = L, \quad b = a/L, \quad i \notin J. \tag{2.1}$$

Thus the positive coordinates have total mass exactly a.

Lemma 2.1 (A block-supported template)**.** *Let m, k, L be positive integers, $a > 0$, $b = a/L$, and $\varepsilon > 0$. Put*

$$Q_0 = \left\lceil \frac{16kL}{\varepsilon} \right\rceil.$$

Suppose $I \subseteq \{1, \dots, m\}$ has at least

$$k + (L+1)Q_0^k$$

elements. If $H \le \mathbb{R}^m$ is a vector subspace of codimension at most k, then there are $i \in I$, $J \subseteq I \setminus \{i\}$ with $|J| = L$, and $h \in H$ such that the template (2.1) satisfies

$$\|h - v\|_{\ell^1(\mathbb{T}^m)} \le \frac{\varepsilon}{16}. \tag{2.2}$$

No rationality assumption on H is needed.

Proof. Choose a coordinate set $C \subseteq \{1, \dots, m\}$ of cardinality $c \le k$ for which

$$H + \mathbb{R}^C = \mathbb{R}^m.$$

Such a set exists because the images of the standard coordinate vectors span $\mathbb{R}^m/H$; select a basis from these images. For every $r \in I \setminus C$, write

$$be_r = h_r + x_r, \qquad h_r \in H, \quad x_r \in \mathbb{R}^C.$$

Reduce the x_r modulo $\mathbb{Z}^C$ and partition the resulting c-torus into Q_0^c half-open boxes of side $1/Q_0$. There are at least $(L+1)Q_0^k$ indices available. Since $c \le k$, one box contains $L+1$ distinct indices. Call one of them i and the others J.

For each $r \in J$, every coordinate of $x_r - x_i$ is, modulo an integer, of absolute value at most $1/Q_0$. Summing gives

$$v = \sum_{r\in J}(be_r - be_i) = \sum_{r\in J}(h_r - h_i) + \sum_{r\in J}(x_r - x_i) = h + x,$$

where $h \in H$, $x \in \mathbb{R}^C$, and

$$\|x\|_{\ell^1(\mathbb{T}^m)} \le \frac{cL}{Q_0} \le \frac{kL}{Q_0} \le \frac{\varepsilon}{16}.$$

This proves the conclusion. In particular, all nonzero coordinates of v lie in the prescribed block I. □

3 The uniform finite nilmanifold input

Here and below a standard presentation is a compact nilmanifold $X = G/\Gamma$ with G connected and simply connected, supplied with a fixed rational Mal'cev presentation and a compatible metric. We enumerate representatives of these presentations as

$$X_1 = G_1/\Gamma_1,\ X_2 = G_2/\Gamma_2, \dots,$$

including the one-point nilmanifold. There are only countably many such presentations up to the isomorphisms needed here. Recurrence is topological, so the choice of one compatible metric for each representative is harmless.

We use normalized polynomial sequences on G, adapted to its lower central series: they have Taylor forms

$$p(n) = g_1^{\binom{n}{1}} \cdots g_s^{\binom{n}{s}}, \qquad g_r \in G_r, \qquad p(0) = 1_G,$$

where G_r denotes the rth lower central subgroup in this paragraph. In particular, $p(n) = g^n$ is allowed.

3.1 The two quoted facts

The nilmanifold facts needed in the construction are the following, in the forms recorded in Green's exposition [2, Lemma 3.1 and Proposition 3.2].

(N1) Every normalized polynomial coset sequence has a representative whose Taylor coefficients belong to fixed compact sets, depending only on the nilmanifold. Evaluation at any fixed integer is continuous as a function of the coefficients.

(N2) On a subprogression $q\mathbb{N}$, a normalized polynomial orbit has closure a connected subnilmanifold through the identity coset. This is the qualitative polynomial orbit-closure/equidistribution theorem of Leibman, as used in Green's Proposition 3.2.

We include the compactness deduction we require, with an arbitrary tolerance rather than a tolerance tied to the number of coordinates.

Lemma 3.1 (Uniform finite orbit approximation)**.** *Let $1, \alpha_1, \dots, \alpha_m$ be linearly independent over $\mathbb{Q}$. Fix a finite collection of standard nilmanifolds and a number $\rho > 0$. There is N such that, for each nilmanifold $X = G/\Gamma$ in the collection and each normalized polynomial sequence p on G, the points*

$$(n\alpha_1, \dots, n\alpha_m, p(n)\Gamma), \qquad 1 \le n \le N,$$

are ρ-dense on some connected subnilmanifold

$$Z = \mathcal{G}/(\mathcal{G} \cap (\mathbb{Z}^m \times \Gamma)) \ \subseteq\ \mathbb{T}^m \times G/\Gamma$$

through the identity coset, with $\dim \mathcal{G} \ge m$. Here distance is the maximum of the sup-coordinate metric on $\mathbb{T}^m$ and the fixed metric on X. The approximating points need not themselves lie in Z, and Z can depend on p.

Proof. It suffices to treat one nilmanifold and then take the maximum of the finitely many bounds. By (N1), restrict attention to a compact parameter space of representatives of normalized polynomial coset sequences.

For one such representative p, apply (N2) to the product sequence

$$n \longmapsto (n\alpha_1, \dots, n\alpha_m, p(n)).$$

On some $q\mathbb{N}$ its closure is a connected subnilmanifold Z_p through the identity. Its projection to $\mathbb{T}^m$ is all of $\mathbb{T}^m$: multiplication by the integer q preserves the rational independence hypothesis, so the torus part on that subprogression is dense. Consequently the presenting subgroup $\mathcal{G}_p$ has dimension at least m.

Compactness of Z_p and density give a finite integer $N(p)$ such that the positive samples through $N(p)$ are $\rho/2$-dense on Z_p. By continuity of these finitely many evaluations, every polynomial sequence in some neighborhood of p has its first $N(p)$ samples ρ-dense on the *same* set Z_p. Notice why we do not require these neighboring samples to lie in Z_p.

These neighborhoods cover the compact parameter space. Take a finite subcover and let N be the maximum of its corresponding $N(p)$'s. Enlarging a finite prefix cannot destroy the required approximation property. This proves the lemma. It is the compactness proof of Green's Proposition 3.3 with ρ in place of $1/m$. □

Remark 3.2 (The available kernel directions)**.** For a subgroup $\mathcal{G}$ in Lemma 3.1, set

$$H \times \{1_G\} = \mathcal{G} \cap (\mathbb{R}^m \times \{1_G\}).$$

Then H is a real vector subspace and

$$\dim H \geq \dim \mathcal{G} - \dim G \geq m - \dim G.$$

One can see this either from the kernel of projection to G or on Lie algebras. Intersections of connected subgroups in a simply connected nilpotent Lie group are connected, by the exponential parametrization. In particular, every $(h \bmod \mathbb{Z}^m, 1_G\Gamma)$ with $h \in H$ belongs to Z. This is the reason that Lemma 2.1 produces a target with nilmanifold coordinate exactly the identity.

4 Constructing one set which retains nilrecurrence

4.1 Parameters fixed once and for all

Set

$$a = \frac{\sqrt{2}}{1000}, \qquad L_j = 10^{12}2^j, \qquad b_j = \frac{a}{L_j}, \qquad \varepsilon_j = \frac{b_j}{100} \quad (j \geq 1). \tag{4.1}$$

In particular, $a/b_j = L_j$ is an integer despite the irrationality of a. We will choose coordinate cutoffs m_j, integer cutoffs N_j, and positive frequencies α_r. Initialize

$$m_0 = 0, \qquad N_0 = 1.$$

Let

$$k_j = \max\{1, \dim G_1, \dots, \dim G_j\}, \qquad Q_j = \left\lceil \frac{16k_jL_j}{\varepsilon_j} \right\rceil,$$

and choose

$$m_j = m_{j-1} + k_j + (L_j + 1)Q_j^{k_j}, \qquad I_j = \{m_{j-1} + 1, \dots, m_j\}. \tag{4.2}$$

The extravagant block sizes are only a convenient way to make the pigeonhole requirements explicit.

4.2 The recursive step

Suppose all frequencies through m_{j-1} and the cutoff N_{j-1} have been fixed. Choose the new frequencies in I_j so that

$$1, \alpha_1, \ldots, \alpha_{m_j} \text{ are linearly independent over } \mathbb{Q}, \qquad \alpha_{r+1} < \alpha_r/2, \tag{4.3}$$

where the latter condition is imposed whenever both sides have been chosen. Also require

$$\alpha_{m_{j-1}+1} < \frac{a}{10N_{j-1}}, \tag{4.4}$$

and, if $j \geq 2$,

$$\alpha_{m_{j-1}+1} < \frac{\varepsilon_{j-1}}{16N_{j-1}}. \tag{4.5}$$

All these requirements are possible: at each choice there is an arbitrarily small positive interval available, and the forbidden rational span of the previously chosen numbers is countable.

Let $\mathcal{V}_j \subset \mathbb{R}^{m_j}$ be the finite collection of all templates

$$v = -ae_i + b_j \sum_{r \in J} e_r, \qquad i \in I_j, \quad J \subset I_j \setminus \{i\}, \quad |J| = L_j. \tag{4.6}$$

Apply Lemma 3.1 to $\alpha_1, \ldots, \alpha_{m_j}$, the first j nilmanifolds, and

$$\rho_j = \min\left\{\frac{1}{2j}, \frac{\varepsilon_j}{16m_j}\right\}. \tag{4.7}$$

Choose $N_j > N_{j-1}$ at least as large as the resulting finite bound. Define

$$D_j = \left\{ n \in \mathbb{N} : \begin{array}{c} N_{j-1} < n \leq N_j, \\ \min\limits_{v \in \mathcal{V}_j} \sum\limits_{r=1}^{m_j} \|n\alpha_r - v_r\|_{\mathbb{T}} < \varepsilon_j/2 \end{array} \right\}. \tag{4.8}$$

This completes stage j. Finally set

$$\boxed{D = \bigcup_{j \geq 1} D_j.} \tag{4.9}$$

No requested density occurs anywhere in these choices.

4.3 Future frequencies do not spoil a completed stage

The geometric decrease gives $\sum_r \alpha_r < \infty$. By the next stage's requirement (4.5),

$$\sum_{r > m_j} \alpha_r < 2\alpha_{m_j+1} < \frac{\varepsilon_j}{8N_j}. \tag{4.10}$$

Thus the map

$$P(n) = (n\alpha_r \bmod 1)_{r \geq 1} \in E$$

is well-defined for every integer n. If $d \in D_j$, choose $v \in \mathcal{V}_j$ as in (4.8) and extend v by zero outside I_j. Then

$$\|P(d) - v\|_E < \frac{\varepsilon_j}{2} + N_j \sum_{r > m_j} \alpha_r < \frac{5\varepsilon_j}{8} < \varepsilon_j. \tag{4.11}$$

Membership in D_j was determined by finite data; all subsequent stages preserve this stronger ambient approximation.

Lemma 4.1 (Consequences for an allowed step)**.** *For $d \in D_j$ there are a distinguished coordinate $i \in I_j$ and centered real representatives σ_r of $d\alpha_r \bmod 1$ such that*

$$\sigma_i = -a + \eta, \qquad |\eta| < \varepsilon_j, \tag{4.12}$$

$$|\sigma_r| \le b_j + \varepsilon_j \qquad (r \ne i), \tag{4.13}$$

$$a - \varepsilon_j < \sum_{r \ne i} |\sigma_r| < a + \varepsilon_j. \tag{4.14}$$

Consequently,

$$\sum_r |\sigma_r| < 2a + 2\varepsilon_j, \qquad \sum_{r \ne i} \sigma_r^2 \le (b_j + \varepsilon_j)(a + \varepsilon_j). \tag{4.15}$$

Proof. The coordinates of v are $-a$, b_j, or 0, all well inside the centered fundamental interval. Since the total torus error in (4.11) is less than ε_j, its coordinate errors can be represented as real errors without a wrap at the ends of that interval. The distinguished error gives (4.12); the individual tail errors give (4.13). The total tail mass differs by less than ε_j from the template's tail mass $L_j b_j = a$. The final two inequalities follow by summation and by bounding the sum of squares by the maximum times the sum. □

4.4 Each stage meets all of its nilrecurrence tests

Proposition 4.2. *For every j, every $X_r = G_r/\Gamma_r$ with $r \le j$, and every normalized polynomial sequence p on G_r, there is $d \in D_j$ such that*

$$d_{X_r}(p(d)\Gamma_r, \Gamma_r) \le \frac{1}{2j}.$$

In particular, every D_j is nonempty.

Proof. Use the subnilmanifold supplied by Lemma 3.1 at stage j. The subspace $H \le \mathbb{R}^{m_j}$ from Remark 3.2 has codimension at most $\dim G_r \le k_j$. By (4.2) and Lemma 2.1, there are $v \in \mathcal{V}_j$ and $h \in H$ with

$$\|h - v\|_{\ell^1(\mathbb{T}^{m_j})} \le \varepsilon_j/16.$$

The point $(h \bmod \mathbb{Z}^{m_j}, \Gamma_r)$ belongs to the subnilmanifold. Therefore some integer $1 \le n \le N_j$ satisfies

$$\max_{u \le m_j} \|n\alpha_u - h_u\|_{\mathbb{T}} \le \rho_j, \qquad d_{X_r}(p(n)\Gamma_r, \Gamma_r) \le \rho_j.$$

The torus estimate and (4.7) give

$$\sum_{u \le m_j} \|n\alpha_u - v_u\|_{\mathbb{T}} \le m_j\rho_j + \varepsilon_j/16 \le \varepsilon_j/8 < \varepsilon_j/2. \tag{4.16}$$

It remains to check that $n > N_{j-1}$; this prevents an old return from being reused at every stage. If $n \le N_{j-1}$, the distinguished coordinate $i \in I_j$ satisfies

$$0 < n\alpha_i < a/10$$

by (4.4). Since $v_i = -a$ and $a + a/10 < 1/2$, it follows that

$$\|n\alpha_i - v_i\|_{\mathbb{T}} = n\alpha_i + a \ge a,$$

contrary to (4.16). Thus $n \in D_j$. The nilmanifold estimate is at most $\rho_j \le 1/(2j)$, as required. The trivial polynomial sequence on any one of the first j nilmanifolds proves nonemptiness. □

Proof of Theorem 1.2(i). First take one standard presentation $X_r = G_r/\Gamma_r$ and $g \in G_r$. For every $j \geq r$, Proposition 4.2 supplies $d_j \in D_j$ with

$$g^{d_j}\Gamma_r \longrightarrow \Gamma_r.$$

The integers d_j tend to infinity because $d_j > N_{j-1}$ and the N_j are strictly increasing.

For an arbitrary point $x = v\Gamma_r$, put $h = v^{-1}gv$. Then

$$g^d v\Gamma_r = vh^d\Gamma_r.$$

Recurrence for $h^d\Gamma_r$ and continuity of left multiplication by v prove recurrence at x.

The standard reduction for nilsystems embeds a given nilsystem in a nilrotation with a connected, simply connected presenting group; see Huang–Shao–Ye [3, Section 3.1.2]. Consequently the same conclusion holds without connectedness assumptions on the original presenting group. Only finiteness of the resulting nilpotency step is needed.

Finally apply pointwise recurrence to the product nilsystem

$$(X^\ell, T \times T^2 \times \cdots \times T^\ell)$$

at $(x, \ldots, x)$. A return to U^ℓ is exactly the simultaneous return asserted in part (i). The preceding construction gives unbounded return times, so the additional condition $d > M$ is automatic after taking a sufficiently late stage. □

5 One fixed obstruction to double topological recurrence

We give the coloring argument directly, so it does not depend on any rationality convenience in the original presentation.

5.1 The complex potential and its second difference

For $n \in \mathbb{Z}$, set

$$\Phi(n) = \sum_{r\geq 1} (\exp(2\pi i n\alpha_r) - 1). \tag{5.1}$$

This is absolutely convergent, since

$$|\exp(2\pi i n\alpha_r) - 1| \leq 2\pi|n|\alpha_r, \qquad \sum_r \alpha_r < \infty.$$

Fix $d \in D_j$ and let the representatives σ_r be those in Lemma 4.1. Then

$$W := \Phi(n) - 2\Phi(n+d) + \Phi(n+2d) = \sum_r \exp(2\pi i n\alpha_r)(1 - \exp(2\pi i\sigma_r))^2. \tag{5.2}$$

The distinguished summand has magnitude

$$|1 - \exp(2\pi i\sigma_i)|^2 = 4\sin^2(\pi\sigma_i).$$

We have $a/2 < |\sigma_i| < 2a < 1/2$. The elementary inequality $\sin(\pi t) \geq 2t$ for $0 \leq t \leq 1/2$ therefore implies that this magnitude is at least $4a^2$.

The remaining summands have total magnitude at most

$$4\pi^2 \sum_{r\neq i} \sigma_r^2 \leq 4\pi^2(b_j + \varepsilon_j)(a + \varepsilon_j) \leq \frac{16\pi^2 a^2}{10^{12}} < a^2. \tag{5.3}$$

For the middle inequality we used $b_j \le a/10^{12}$, $\varepsilon_j \le b_j$, and $a+\varepsilon_j < 2a$. For an upper bound on the distinguished term use $|1-\exp(2\pi i t)| \le 2\pi|t|$. Together these estimates give, uniformly in n, d, j,

$$\boxed{3a^2 < |W| < (16\pi^2+1)a^2 < \frac{1}{4}.} \tag{5.4}$$

The last inequality follows from $a^2 = 2\cdot 10^{-6}$.

5.2 A finite coloring

Partition the half-open square $[0,1)^2 \cong \mathbb{C}/\mathbb{Z}[i]$ into finitely many half-open squares of Euclidean diameter at most

$$\zeta = a^2/100.$$

Color $n \in \mathbb{Z}$ according to the square containing $\Phi(n)$ mod $\mathbb{Z}[i]$. Suppose $n, n+d, n+2d$ had the same color for some $d \in D$. Write

$$\Phi(n+td) = \lambda_t + r_t, \qquad \lambda_t \in \mathbb{Z}[i], \quad t = 0,1,2,$$

where the r_t lie in one of the chosen squares. Then

$$W = \Lambda + R, \qquad \Lambda = \lambda_0 - 2\lambda_1 + \lambda_2 \in \mathbb{Z}[i], \qquad R = (r_0 - r_1) + (r_2 - r_1),$$

and $|R| \le 2\zeta$. Since $|W| < 1/4$ and $2\zeta < 1/4$, a nonzero Gaussian integer Λ is impossible. Hence $\Lambda = 0$, giving

$$|W| \le 2\zeta = a^2/50,$$

contrary to (5.4). This proves the claimed absence of monochromatic three-term progressions.

5.3 The corresponding minimal dynamical witness

Let $c \in \{1,\dots,r_0\}^{\mathbb{Z}}$ be the coloring just constructed and let Y_0 be its orbit closure under the shift. For every $d \in D$, no point of Y_0 has equal symbols at coordinates $0, d, 2d$. Indeed the same is true for every shift of c, and the condition that these three symbols are equal describes a clopen cylinder.

Choose a nonempty minimal subsystem $Y \subseteq Y_0$ and a symbol e for which

$$V = \{y \in Y : y_0 = e\}$$

is nonempty. This set is clopen, and

$$V \cap S^{-d}V \cap S^{-2d}V = \varnothing \qquad (d \in D).$$

This proves Theorem 1.2(ii), including the usual minimal-open-set formulation of failure of double topological recurrence.

6 Vanishing tail curvature and narrow forbidden bands

We next prove the analytic/combinatorial part of the density obstruction. It works for every real translate of the potential, which will be important for averaging.

6.1 A discrete estimate

For a real sequence write $\Delta z_t = z_{t+1} - z_t$ and $\Delta^2 z_t = z_{t+2} - 2z_{t+1} + z_t$.

Lemma 6.1 (Drift from range and curvature)**.** *Suppose $z_0, \dots, z_{K-1}$ has oscillation at most M, and $|\Delta^2 z_t| \le B$ whenever defined. If $Q \ge 1$ and $t + Q \le K - 1$, then*

$$|\Delta z_t| \le \frac{M}{Q} + \frac{B(Q-1)}{2}.$$

Proof. Summing first differences, and expressing each in terms of the first one, gives the exact identity

$$z_{t+Q} - z_t = Q\Delta z_t + \sum_{r=0}^{Q-2} (Q - 1 - r)\Delta^2 z_{t+r}.$$

The left side has magnitude at most M, and the last sum has magnitude at most $BQ(Q-1)/2$. Divide by Q. The formula also holds for $Q = 1$, when the sum is empty. □

6.2 Two periodic barriers

Define the real potential

$$F(n) = \Re\Phi(n) = \sum_r \big(\cos(2\pi n\alpha_r) - 1\big). \tag{6.1}$$

For an integer $P \ge 2$ and $0 < h < 1/4$, let

$$\mathcal{R}_{P,h} = \{y \in \mathbb{R} : y \bmod P \in [1, P),\ y \bmod 1 \in [h, 1)\}. \tag{6.2}$$

The coarse condition deletes intervals of width 1 every P units. The fine condition deletes intervals of width h every unit. In one period $[0, P)$ their common retained set is exactly

$$\bigcup_{k=1}^{P-1} [k + h, k + 1),$$

so its normalized length is

$$\beta = \frac{(P-1)(1-h)}{P} = (1 - 1/P)(1 - h). \tag{6.3}$$

For $\theta \in \mathbb{R}$ define

$$B_\theta = \{n \in \mathbb{N} : F(n) + \theta \in \mathcal{R}_{P,h}\}.$$

Proposition 6.2 (Every translate obstructs all sufficiently late stages)**.** *For every $P \ge 2$ and $0 < h < 1/4$, there are integers $J \ge 2$ and $K_0 \ge 3$ such that, for every $\theta \in \mathbb{R}$, the set B_θ contains no K_0-term arithmetic progression with difference in $\bigcup_{j \ge J} D_j$.*

Proof. We specify the constants before examining a progression.

Step 1: choose a uniform sampling length. Since a is irrational, there is an integer $L \ge 2$ for which

$$\{0, a, \dots, (L-1)a\} \subseteq \mathbb{T}$$

is $h/(16\pi)$-dense. Every translate and reflection of this finite set has the same property. Thus, for every real c and every phase x, among the values

$$c + \cos(2\pi(x - ta)) - 1, \qquad 0 \le t < L, \tag{6.4}$$

there is one at distance at most $h/8$ from $\mathbb{Z}+h/2$. To justify this, the continuous range of $c+\cos(2\pi s)-1$ is $[c-2,c]$, which contains a point of $\mathbb{Z}+h/2$. Choose a phase realizing that value and approximate it by some $x-ta$. The cosine is 2π-Lipschitz in the circle metric.

Step 2: choose the confinement length and then the stage. Choose an integer

$$Q \geq \frac{16L(P+1)}{h}, \qquad K_0 = Q + L. \tag{6.5}$$

Put

$$B_j = 4\pi^2(b_j+\varepsilon_j)(a+\varepsilon_j).$$

We have $B_j \to 0$ and $\varepsilon_j \to 0$. Choose $J \geq 2$ so large that, for every $j \geq J$,

$$\frac{B_j(Q-1)}{2} \leq \frac{h}{16L}, \qquad 2\pi L\varepsilon_j \leq \frac{h}{8}. \tag{6.6}$$

These choices are independent of a starting integer, an allowed difference, and θ.

Step 3: coarse confinement on a hypothetical progression. Suppose

$$n, n+d, \dots, n+(K_0-1)d \in B_\theta, \qquad d \in D_j, \quad j \geq J.$$

Let i, σ_r be given by Lemma 4.1, and set

$$y_t = F(n+td)+\theta, \qquad u_t = \cos(2\pi(n\alpha_i + t\sigma_i))-1, \qquad z_t = y_t - u_t.$$

By (4.15),

$$|y_{t+1}-y_t| \leq 2\pi\sum_r |\sigma_r| < 2\pi(2a+2\varepsilon_j) < 1. \tag{6.7}$$

Every y_t satisfies the coarse retained condition. Two distinct coarse components are separated by a gap of width 1, so (6.7) prevents the sequence from changing components. Therefore

$$\operatorname{osc} y_t \leq P-1, \qquad \operatorname{osc} z_t \leq (P-1)+2 = P+1.$$

The last bound uses only that a single cosine minus one has range $[-2,0]$.

Step 4: the tail is nearly constant on the first L samples. For a single cosine, its second difference at step v has absolute value at most $4\pi^2v^2$. Summing over the nondistinguished coordinates gives

$$|\Delta^2 z_t| \leq 4\pi^2 \sum_{r\neq i} \sigma_r^2 \leq B_j. \tag{6.8}$$

All sums are justified by absolute convergence; alternatively the estimates follow first for finite truncations and then by passing to the limit.

For $0 \leq t \leq L-1$ we have $t+Q \leq K_0-1$. Apply Lemma 6.1 with $M = P+1$ and use (6.5)–(6.6). We obtain

$$|\Delta z_t| \leq \frac{P+1}{Q} + \frac{B_j(Q-1)}{2} \leq \frac{h}{8L}.$$

In particular,

$$|z_t - z_0| \leq h/8 \qquad (0 \leq t < L). \tag{6.9}$$

Step 5: an ideal irrational cosine forces a fine-band hit. Write $\sigma_i = -a+\eta$, with $|\eta| < \varepsilon_j$. For $0 \leq t < L$, the Lipschitz bound and (6.6) give

$$\bigl|u_t - (\cos(2\pi(n\alpha_i - ta))-1)\bigr| \leq 2\pi t\varepsilon_j \leq h/8.$$

Combining with (6.9),

$$\big|y_t - \big(z_0 + \cos(2\pi(n\alpha_i - ta)) - 1\big)\big| \le h/4 \qquad (0 \le t < L). \tag{6.10}$$

Step 1 supplies some $t < L$ for which the ideal value in (6.10) is at distance at most $h/8$ from an integer plus $h/2$. The actual value y_t is therefore at distance at most $3h/8$ from that center. Hence

$$y_t \bmod 1 \in [h/8, 7h/8] \subset (0, h),$$

contradicting the fine retained condition in (6.2). This proves the proposition for every θ. □

Remark 6.3 (Where the changes are used)**.** Irrationality of the *fixed* a is used in Step 1, for arbitrarily small h. The limits $b_j, \varepsilon_j \to 0$ are used only after P, h, L, Q have been fixed, in Step 2. The old fixed-parameter construction does not automatically supply either of these two features. No iteration of a fixed-density estimate is being assumed.

7 All densities: phase averaging, early stages, and exact thinning

We complete Theorem 1.2(iii).

7.1 Choose the real barriers

Fix $0 < \delta < 1$. Choose an integer $P \ge 2$ and $0 < h < 1/4$ so that

$$\beta = (1 - 1/P)(1 - h) > \delta.$$

For example, choosing both $1/P$ and h smaller than $(1-\delta)/4$ suffices. Take J, K_0 from Proposition 6.2. The initial exceptional set

$$E_0 = \bigcup_{1 \le j < J} D_j$$

is finite and lies in $[1, N_{J-1}]$.

7.2 An arbitrarily cheap periodic deletion handles all early differences

Choose a prime $q > N_{J-1}$ so large that

$$\beta(1 - 1/q) > \delta. \tag{7.1}$$

For $\theta \in [0, P)$ put

$$A_\theta = \{n \in \mathbb{N} : q \nmid n,\ F(n) + \theta \in \mathcal{R}_{P,h}\}. \tag{7.2}$$

For $d \in E_0$, we have $0 < d < q$, and thus $\gcd(d, q) = 1$. Any q consecutive terms of an arithmetic progression with difference d run through all residues modulo q. One is divisible by q, so such a progression cannot lie in A_θ.

For $d \in \bigcup_{j \ge J} D_j$, the set $A_\theta \subseteq B_\theta$ contains no K_0-term progression with difference d. Consequently, for every θ,

$$A_\theta \text{ contains no } K\text{-term progression with difference in } D, \qquad K = \max\{q, K_0, 3\}. \tag{7.3}$$

Notice that increasing q reduces the density cost of handling the early stages; it does not modify D.

7.3 Phase averaging gives the required upper density

No distribution theorem for the infinite sum F is needed. A real translate of a P-periodic set of relative measure β has average indicator β. Therefore, for every N,

$$\begin{aligned}\frac{1}{P}\int_0^P \frac{|A_\theta \cap [1,N]|}{N}\,d\theta &= \frac{1}{N}\sum_{\substack{1\le n\le N\\ q\nmid n}} \frac{1}{P}\int_0^P \mathbf{1}_{\mathcal{R}_{P,h}}(F(n)+\theta)\,d\theta \\ &= \beta\left(1-\frac{\lfloor N/q\rfloor}{N}\right). \end{aligned} \tag{7.4}$$

The functions being averaged lie between zero and one. Reverse Fatou's inequality, equivalently Fatou applied to their complements, implies

$$\frac{1}{P}\int_0^P \overline{d}(A_\theta)\,d\theta \ge \limsup_{N\to\infty}\frac{1}{P}\int_0^P \frac{|A_\theta\cap[1,N]|}{N}\,d\theta = \beta(1-1/q) > \delta. \tag{7.5}$$

Thus there exists a *single* phase θ such that

$$\alpha := \overline{d}(A_\theta) > \delta.$$

The phase does not vary with N. Measurability is immediate from (7.2): each finite counting function is measurable, and a limsup of measurable functions is measurable.

7.4 Thinning to upper density exactly δ

Lemma 7.1 (Exact upper-density thinning)**.** *If $A\subseteq\mathbb{N}$ has $\overline{d}(A)=\alpha>0$ and $0<\delta\le\alpha$, then there is $A'\subseteq A$ with $\overline{d}(A')=\delta$.*

Proof. Enumerate $A=\{a_1<a_2<\cdots\}$ and put $\lambda=\delta/\alpha\le 1$. Retain a_j exactly when

$$\lfloor \lambda j\rfloor > \lfloor \lambda(j-1)\rfloor.$$

The resulting set satisfies the exact counting identity

$$|A'\cap[1,N]| = \lfloor \lambda |A\cap[1,N]|\rfloor.$$

After division by N and taking limsups, this gives $\overline{d}(A')=\lambda\alpha=\delta$. □

Apply the lemma to the phase selected in (7.5), and call the resulting subset A_δ. Since a subset cannot create an arithmetic progression which was absent before thinning, (7.3) holds for A_δ. Hence $\overline{d}(A_\delta)=\delta$ and A_δ has no K-term progression with difference in D. This completes the proof of Theorem 1.2.

8 Why the diagonalization is legitimate

We summarize the dependencies to make two possible pitfalls explicit.

The construction is not an arbitrary intersection of nilrecurrent sets. At stage j, one first chooses a new coordinate block large enough for every codimension up to k_j, then its independent frequencies, and only then a cutoff N_j which works uniformly for all normalized polynomial sequences on the first j nilmanifolds. Every one of these tests has a witness in the finite set D_j. The new block's small frequencies ensure that this witness occurs beyond N_{j-1}. Future frequencies are constrained by a summable tail budget, so they preserve every earlier membership assertion.

The density is not a construction parameter. The data $a, L_j, b_j, \varepsilon_j, m_j, \alpha_r, N_j, D_j, D$ are all fixed before any density is given. A requested density determines P, h, L, Q, J, then a prime q, then a phase θ, and finally a subset selected by rank. Shrinking $1-\delta$ merely moves the useful tail estimate to later stages and increases the forbidden progression length.

Both recurrence conclusions concern the same set. Every stage retains nilrecurrence tests while obeying the same complex annulus estimate (5.4). Thus taking the union of the stages preserves a *single* finite-color obstruction to double topological recurrence. The density argument handles all sufficiently late stages uniformly, with one periodic deletion taking care of the finitely many initial stages.

External inputs. The nilmanifold inputs are the compact coefficient reduction and polynomial orbit closure stated in Section 3, plus the standard connected, simply connected reduction when allowing general nilsystem presentations. All subsequent approximation, diagonalization, coloring, and density arguments are given above. The result establishes the requested modified example; it makes no claim about the still separate all-densities question for the unchanged example.

References

[1] R. Alweiss, *New obstacles to multiple recurrence*, arXiv:2511.21680, version 3, 2025. https://arxiv.org/abs/2511.21680.

[2] B. Green, *On Alweiss's example for multiple recurrence*, arXiv:2607.24594, version 1, 2026. In particular Lemma 3.1, Proposition 3.2, and the compactness proof of Proposition 3.3. https://arxiv.org/abs/2607.24594.

[3] W. Huang, S. Shao, and X. Ye, *Nil Bohr-sets and almost automorphy of higher order*, arXiv:1407.1179, version 1, 2014. Section 3.1.2 gives the connected, simply connected reduction used here. https://arxiv.org/abs/1407.1179.