Throughout, $\pi = \pi^{[n]}$ denotes a fixed product distribution on $[k]^{[n]}$.

Theorem 1 *Let $B \subseteq [k]^n$ be ab-insensitive on I, and assume $|I| \geq n_{\text{tile}}(k, d, \eta, d) := 2m(k + d)^m \ln(1/\eta)$, where $m = n_{\text{subsp}}(k, d, \eta/2, \pi)$.*[1] *Then B can be partitioned into a collection $\mathcal{S}$ of disjoint nondegenerate d-dimensional subspaces, along with an error set E satisfying $\pi(E) \leq \eta$.*

Proof: The proof proceeds in "rounds", $t = 1, 2, 3, \ldots$. In each round, we remove some disjoint subspaces from B and put them into $\mathcal{S}$; and, the set I shrinks by m coordinates. We will show that after the tth round,

$$\pi(B) \leq \left(1 - \frac{1}{2(k + d)^m}\right)^t. \tag{1}$$

We continue the process until $\pi(B) \leq \eta$, at which point we may stop and set $E = B$. Because of (1), the process stops after at most $T = 2(k + d)^m \ln(1/\eta)$ rounds. By our choice of $n_{\text{tile}}(k, d, \eta, d) = mT$, the set I never "runs out of coordinates".

Suppose we are about to begin the tth round; hence, writing $\delta = \pi(B)$ we have

$$\eta < \delta \leq \left(1 - \frac{1}{2(k + d)^m}\right)^{t-1}.$$

The round begins by choosing an arbitrary $J \subseteq I$ with $|J| = m$. Since π is a product distribution we have

$$\delta = \Pr_{x \sim \pi^{[n]}}[x \in B] = \mathop{\mathbf{E}}_{y \sim \pi^{[n] \setminus J}}[\pi^J(B_y)],$$

where $B_y = \{z \in [k]^J : (y, z) \in B\}$. Hence $\pi^J(B_y) \geq \delta/2$ for at least a $\delta/2$ probability mass of y's; call these y's "good". Since $|J| = m = n_{\text{subsp}}(k, d, \eta/2, \pi) \geq n_{\text{subsp}}(k, d, \delta/2, \pi)$, for each good y there must exist a nondegenerate d-dimensional subspace $\rho \subseteq B_y$. Since the number of d-dimensional subspaces in $[k]^m$ is at most $(k + d)^m$, there must exist a *fixed* nondegenerate d-dimensional subspace $\rho_0 \subseteq [k]^J$ such that

$$\Pr_{y \sim \pi^{[n] \setminus J}}[\rho_0 \subseteq B_y] \geq \frac{\delta}{2(k + d)^m}. \tag{2}$$

Let $R \subseteq [n] \setminus J$ be the set of y's with $\sigma_0 \subseteq B_y$. It is easy to see that R is ab-insensitive on $I \setminus J$. Hence $R \times \rho_0$ is ab-insensitive on $I \setminus J$; hence so too is $B \setminus (R \times \rho_0)$. We therefore complete the round by setting $I = I \setminus J$ and transferring $R \times \rho_0$ (a disjoint union of subspaces $\{y\} \times \rho_0$) from B into $\mathcal{S}$. By (2), this drops B's probability mass from δ to $\delta \cdot (1 - 1/2(k + d)^m)$, as required to establish (1) inductively. $\square$

[1] I.e., m is large enough that every ab-insensitive set $B \subseteq [k]^m$ with $\pi(B) \geq \eta/2$ contains a nondegenerate d-dimensional subspace.