HW1

Mathematical Methods in TCS: Boolean Functions (Fall 2026)
University of Toronto
Swastik Kopparty

Due: Wednesday, Dec 6, 2026

Answer all questions.

Do not use AI to give you answers to these questions.

1. Let V be an $\mathbb{F}_2$-linear subspace of $\mathbb{F}_2^n$. Define $V^\perp$ to be its dual subspace:

$$V^\perp = \{y \mid \langle x, y\rangle = 0 \forall x \in V\}.$$

   Express $\widehat{1_V}$ in terms of $1_{V^\perp}$.

2. We have an $\mathbb{F}_2^n$-valued random variable $(X_1, \ldots, X_n)$. (The X_i are thus $\{0,1\}$-valued, and need not be independent.)

   Let $\mu : \mathbb{F}_2^n \to \mathbb{R}$ be the probability mass function of the distribution of $(X_1, \ldots, X_n)$.

   The distribution of $(X_1, \ldots, X_n)$ is called k-wise independent if (1) each X_i is uniformly distributed in $\{0,1\}$, and (2) for any $S \subseteq [n]$ with $|S| \leq k$, the random variables $(X_j)_{j \in S}$ are independent.

   (a) Show that μ is k-wise independent if and only if $\hat{\mu}(S) = 0$ for all $S \subseteq [n]$ with $1 \leq |S| \leq k$.

   (b) Let k be a constant. Let $X \subseteq \mathbb{F}_2^n$ with $|X| = o(n^{\lfloor k/2 \rfloor})$. Show that for $d = \lfloor k/2 \rfloor$, there is a function $f : \mathbb{F}_2^n \to \mathbb{R}$ which satisfies:

      i. f is not identically 0.
      ii. $\hat{f}(S) = 0$ for each $|S| > d$.
      iii. $f(x) = 0$ for each $x \in X$.

      Consider the inner product $\langle f, \mu\rangle$. Use it to show that any k-wise independent distribution over $\mathbb{F}_2^n$ has support at least $\Omega(n^{k/2})$.

      (This is tight. There is a clever construction of a small-support k-wise independent distribution using "BCH codes".)

3. We have an $\mathbb{F}_2^n$-valued random variable $(X_1, \ldots, X_n)$. (The X_i are thus $\{0,1\}$-valued, and need not be independent.)

   Let $\mu : \mathbb{F}_2^n \to \mathbb{R}$ be the probability mass function of the distribution of $(X_1, \ldots, X_n)$.

   The distribution of $(X_1, \ldots, X_n)$ is called ϵ-biased if for every nonempty $S \subseteq [n]$:

$$\Pr[\sum_{i \in S} X_i = 1] \in [\frac{1-\epsilon}{2}, \frac{1+\epsilon}{2}].$$

(a) Show that $(X_1, \ldots X_n)$ is ϵ-biased if and only if $|\hat{\mu}(S)| \leq \frac{1}{2^n} \cdot \epsilon$ for each $S \neq \emptyset$.

(b) Show that for any ϵ-biased distribution μ, $\sum_{x \in \mathbb{F}_2^n} \mu(x)^2 \leq \epsilon^2 + \frac{1}{2^n}$. Conclude that if $\epsilon < \frac{1}{2^{n/2}}$, then μ must have support at least $\frac{1}{2}2^n$.

(c) Let μ be an ϵ-biased distribtion. Let t be a parameter to be chosen later. Let $\nu = \mu \star \mu \star \ldots \star \mu$ (t times).

i. Show that ν is an ϵ^t-biased distribution.

ii. Show that $|\mathsf{Support}(\nu)| \leq \binom{|\mathsf{Support}(\mu)|+t}{t}$.

Combine these with the previous part to show that

$$|\mathsf{Support}(\mu)| \geq \Omega(\frac{n}{\epsilon^2 \log \frac{1}{\epsilon}}).$$

This theorem and its proof are due to Noga Alon. Alon's proof is actually even simpler, using only ranks and eigenvalues of matrices, and does not use any Fourier analysis. The original paper (Perturbed Identity Matrices have High Rank) is highly recommended.