1 Learning low-degree $\mathbb{F}_2$ polynomials

Lemma 1.1 *Let $\boldsymbol{X}$ be a (multi)set of $m := 2^e \cdot O(\log(2^{n^e}/\delta)) = n^e \cdot O(2^e \log(1/\delta))$ points drawn uniformly and independently from $\mathbb{F}_2^n$. Then except with probability at most δ it holds that for all nonzero $\mathbb{F}_2$-multilinear polynomials q of degree at most e, $q(x) \neq 0$ for at least one $x \in \boldsymbol{X}$.*

Proof: By a problem from Homework #2 (modified for $\mathbb{F}_2$-multilinear polynomials as opposed to $\mathbb{R}$-multilinear polynomials), for any *specific* q of degree at most e, $\mathbf{Pr}_{\boldsymbol{x}\in\mathbb{F}_2^n}[q(\boldsymbol{x}) \neq 0] \geq 2^{-e}$. Hence

$$\mathop{\mathbf{Pr}}_{\boldsymbol{X}}[q(x) = 0 \;\; \forall x \in \boldsymbol{X}] \leq (1 - 2^{-e})^m \leq \delta/2^{n^e} \leq \delta/(\text{\# degree} \leq e \text{ polys}).$$

The result now follows from the union bound. □

Lemma 1.2 *Let $\boldsymbol{X}$ be a (multi)set of the same number of* examples, $(x, f(x))$, *where x is drawn uniformly from $\mathbb{F}_2^n$ and f is expressible as some $\mathbb{F}_2$-multilinear polynomial of degree at most e. Then except with probability at most δ over the choice of $\boldsymbol{X}$, there is only one polynomial p of degree at most e consistent with the data $\boldsymbol{X}$, namely f.*

Proof: Otherwise, if $p \not\equiv p'$ are both consistent with $\boldsymbol{X}$, then $q := p - p'$ is a nonzero polynomial of degree at most e which is 0 on all the points in $\boldsymbol{X}$. The result follows from the previous lemma. □