

abstract

COMPUTER SCIENCE/DISCRETE MATH I

Topic:

Speaker:

Affiliation:

Date:

Time/Room:

A halfspace (a.k.a. threshold) is a boolean function $h : \{-1,1\}^n \rightarrow \{-1,1\}$ of the form $h(x) = \text{sign}(w_0 + w_1 x_1 + \dots + w_n x_n)$, where the weights w_i are arbitrary real numbers. Halfspaces are studied extensively in the theories of complexity, learning, and social choice.

In this work we show that every distribution on $\{-1,1\}^n$ that is k -wise independent fools any halfspace with error ϵ for $k < (1/\epsilon)^{2+o(1)}$. Our result matches up to polylogarithmic factors the lower bound $k > (1/\epsilon)^{2-o(1)}$ by Benjamini, Gurel-Gurevich, and Peled (2007). A corollary of our result is an explicit pseudorandom generator $G : \{-1,1\}^s \rightarrow \{-1,1\}^n$ that fools halfspaces with error ϵ and seed length $s = (\log n)(1/\epsilon)^{2+o(1)}$, which for constant ϵ is optimal up to constant factors. This is the first generator for halfspaces.

Our approach combines classical tools from real approximation theory with structural results on halfspaces by Servedio (Computational Complexity, 2007).

Joint work with Ilias Diakonikolas, Parikshit Gopalan, Ragesh Jaiswal, and Rocco Servedio.