Title: A Nearly Optimal Solution for the Chow Parameters Problem and Applications
Speaker: Ilias Diakonikolas

Abstract: The Chow parameters of an n-variable Boolean function are its (n + 1) degree-0 and degree-1 Fourier coefficients. It has been known since 1961 that the (exact values of the) Chow parameters of any linear threshold function (LTF) f uniquely specify f within the space of all Boolean functions, but until recently nothing was known about efficient algorithms for reconstructing f (exactly or approximately) from exact or approximate values of its Chow parameters. We refer to this reconstruction problem as the Chow Parameters Problem.

In this talk I will describe a new algorithm for the Chow Parameters Problem which, given (sufficiently accurate approximations to) the Chow parameters of any LTF f, runs in time O(n^2)* (1/eps)^O(log^2 (1/eps)) and outputs a representation of an LTF g that is eps-close to f . The only previous algorithm [O'Donnell - Servedio, STOC'08/SICOMP'11] had running time poly(n)*2^{2^{O(1/eps^2)}}. As a byproduct of our approach, we obtain nearly optimal low-weight approximations for LTFs and improved algorithms for several problems in learning theory.

The two main ingredients underlying our results are (1) a new structural result showing that for f any LTF and g any bounded function, if the Chow parameters of f are close to the Chow parameters of g, then f is close to g (in L_1 distance); (2) a new boosting-like algorithm that given approximations to the Chow parameters of an LTF outputs a bounded function whose Chow parameters are close to those of f.

The talk will be based on joint work with Anindya De (Berkeley), Vitaly Feldman (IBM Almaden), and Rocco Servedio (Columbia).