

**The University of Chicago**

Department of Statistics

**Seminar**

---

**David McAllester**

ATT Labs-Research

**“The PAC-Bayesian Approach to Generalization Guarantees”**

\*\*\*\*\*

**Monday, November 25, 2002 at 4:00 PM**  
**133 Eckhart Hall, 5734 S. University Avenue**

**ABSTRACT**

The PAC-Bayesian theorem provides a generalization guarantee for a Gibbs (stochastic) classifier in terms a prior and a posterior distribution on a set of hypotheses. The performance guarantee is in the PAC tradition (Probably Approximately Correct) and is "nonBayesian" in the sense that there is no assumption that data is generated in any way from the prior. However, the theorem is "Bayesian" in the sense that it applies to an arbitrary prior and specifies an optimal posterior given a prior and training data. The PAC-Bayesian theorem provides an alternative to VC theory in proving a variety of performance guarantees. Applications to linear threshold classifiers and Gaussian processes will be discussed in detail.

---

10/9/02