

COLLOQUIUM

WESLEY CALVERT

Department of Mathematics,
SOUTHERN ILLINOIS UNIVERSITY

9-17-15

NECKERS 156 | 3PM

RECEPTION IMMEDIATELY FOLLOWING IN THE MATH LIBRARY

MATHEMATICS OF THEORETICAL MACHINE LEARNING

[ABSTRACT]

What can a computer learn? Since the 1960s, there have been mathematical models of this problem. The usual modern formulation allows for the computer to be given a sample of data, with each point labeled correct or incorrect. The computer is then asked to identify a rule which will, with high probability, be correct enough to decide whether a new point is correct or incorrect, with perhaps some small region of error. Some cases where this is possible are well-known. Other cases where this is impossible are also well-known. Other (even more well-known) cases are things that people try, and then let the market sort out whether it works well enough (think of Netflix, trying to learn what movies you'd enjoy).

The central problem of this talk is the difficulty of deciding whether something is learnable or not. Logicians have many tools for thinking about the "difficulty" of various problems, and I'll introduce some of them. Much of the real show, though, is in finding the right representations of the problem.