

**University of California, Irvine
Statistics Seminar**

Deep Learning: A Synthesis from Probabilistic Foundations

**Padhraic Smyth, Professor
Eric Nalisnick, PhD
Department of Computer Science, UCI**

**Wednesday, May 9, 2018
4 p.m., 6011 Bren Hall
(Bldg. #314 on campus map)**

Deep Learning (DL)—a family of neural-network-based techniques—has received widespread attention for its impressive performance on a range of pattern recognition problems. Over the past five years, DL systems have significantly improved upon benchmarks in computer vision, reinforcement learning, and speech processing, to name just a few application areas. However, DL research is conducted mostly within the machine learning community, and as a by-product of this isolation, some methodologies have become detached from their heritage in probabilistic modeling. In this talk, we aim to give a statistician-centric view of DL. We review the three primary DL architectures—feedforward networks, recurrent networks, and autoencoders—and show their relationship to models well-known to statisticians, such as generalized linear models, state-space models, and factor analysis (respectively). We will also discuss several broad themes that guide DL research and leave the audience with open problems particularly pertinent to the expertise of statisticians.