Features and Representations

Misha Belkin
Ohio State University
Department of Computer Science and Enginnering,
Department of Statistics

Features and architecture

- Can features be learned?
- Deep or shallow structures?
- Model the architecture or the output?

An example from learning Boolean circuits.

Learning Boolean circuits

Problem: learn a constant depth Boolean circuit.

Possibility: search over the possible architectures.
[Let me know when you are done].

Solution: Fourier analysis on a binary cube. [Linial, Mansour, Nisan, JACM 93]

Fit a bandlimited function (small number of harmonics) to the output of the circuit.

$$f(x) = \sum_{i=1}^{k} a_i e_i(x_i)$$

Sample on random points, then fit.

Equivalently dimensionality reduction..: linear regression in Fourier features.

Aspects of the solution

1. Appropriate (Fourier) features make the algorithm (but not the analysis!) trivial (linear regression).

2. Learn the output, not the architecture.

3. The internal principles very different from the apparent structure.

Applies?