CME 216, ME 343 - Spring 2020

Eric Darve, ICME

Machine learning in computational engineering

Computational engineering has unique requirements, which are different from image and signal processing, or computer vision.

Regression is a much more common application than classification.

In most cases, we are interested in computing a pressure, velocity, or strain field.

The required accuracy is typically high.

For example, we often expect errors in the range 10^{-4}–10^{-3} or less to have confidence in the prediction.

Data is more limited

Data in computational engineering is typically obtained either through

- experiments, which may be expensive and typically do not generate large amounts of data, or
- high-fidelity computer simulations.

The advent of automated sensors and large-scale parallel computers is changing this landscape.

We are often interested in quantifying the uncertainty in the prediction and deriving intervals of confidence for our prediction.

We can leverage equations from physics, conservation laws, physical principles, which can inform our learning algorithms.

This is different from other applications in computer science for example where most of the learning is restricted to learning from data.

See some recent research papers on this topic.