SAND2017-5192D

Adaptive Learning Theory

Sandia National Laboratories
Craig M. Vineyard, Ojas Parekh, Cindy Phillips, James B. Aimone, and Conrad James

Problem

- One of the differentiating capabiltiies of the brain is continuous learning
- However, most data-driven algorithms in ML do not continuously adapt

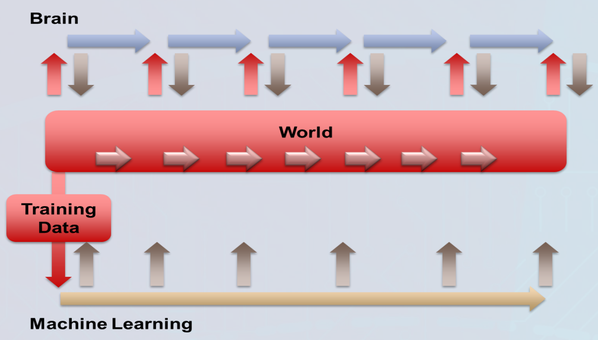

- When should models be re-trained or adapted?

Lots of Ways the World Can Change

- Sample data change patterns over time

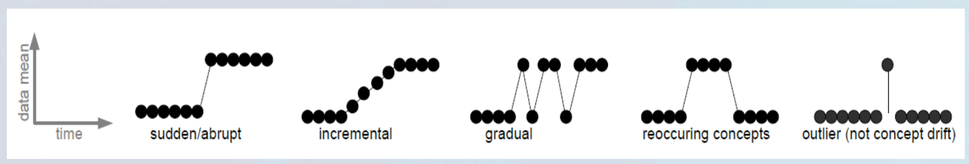

Gama, João, et al. "A survey on concept drift adaptation." *ACM Computing Surveys (CSUR)* 46.4 (2014): 44.

Motivation

- Machine learning algorithms do not have a lack of learning paradigms. In fact there are many
  - For instance supervised and semi-supervised paradigms address how to handle labeled data
  - Methodologies such as batch, incremental, one-shot, and online address how data is presented to learning algorithms
- ...but they have limitations

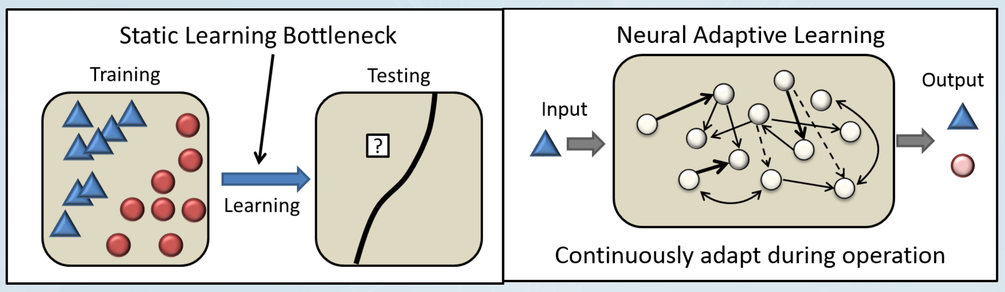

Vineyard, C.M. et al. "Overcoming the Static Learning Bottleneck – the Need for Adaptive Neural Learning." *ICRC 2016*.

Static Learning Bottleneck - distinct training and testing phases necessitate that for a model to be updated and learn it must be re-trained before it may be employed

Cyber Defense Example

- Regex Golf aims to create the smallest regular expression set which accepts one list and rejects another list
  - In cyber security, as threats change (such as concept drift), some indicators are no longer useful at tracking a threat and new indicators/new threats may be added

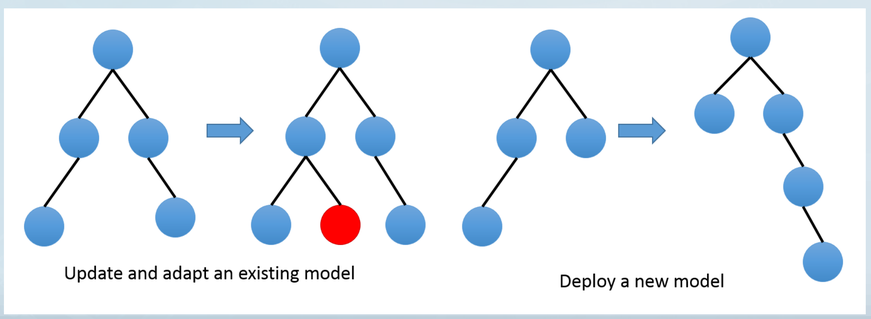

Approach

PLADD Overview

- Probabilistic Learning Attacker, Dynamic Defender (PLADD)
- Game theoretic model and analysis of moving target defense (MTD)
- MTD asserts the use of randomization, diversity, or change make a computer system more difficult to attack (make it a "moving target")

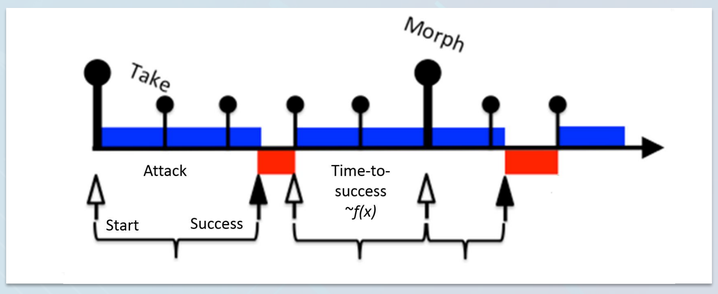

- Two players: attacker and defender
- One contested resource. Defender holds at start
- A player can move at a cost
  - The "take" move - seizes control of the resource immediately
  - The "morph" move - resets the game
  - Neither player ever knows who owns the resource
- Strategy: when to move? Timeline is infinite.
- Utility = (time in control) – cost (can be weighted)

Adaptive Learning Game

- Game theoretic model and analysis of benefits of various forms of adaptation in learning

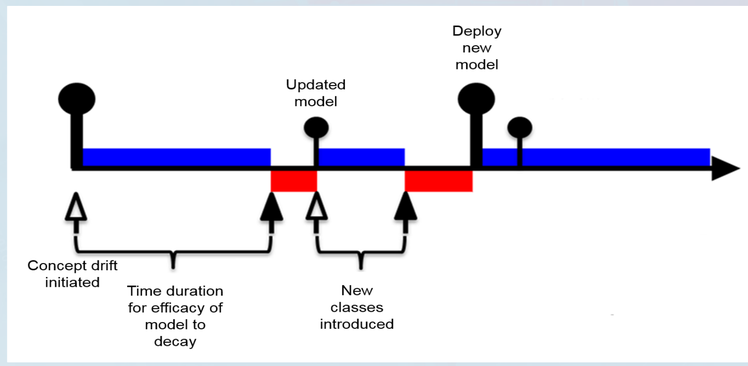

Encoding Theory

- Fundamental to learning theory is a representation of some sort (functional, encoding, etc.) which is manipulated
- Quantifiable traits of an encoding provide insight into the behavior of algorithms
  - Optimal Binary Search Trees
  - Combinatorial Scheduling Problem

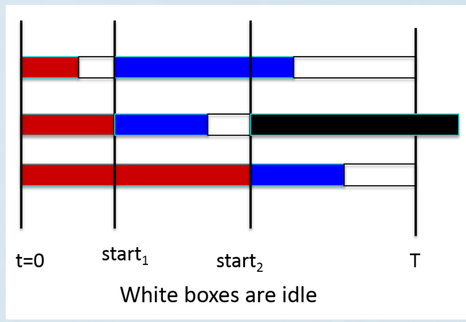

- *m* machines – each corresponds to a scenario
- Each machine has a set of jobs that must be run in order
- Problem: Schedule *t* "global start times" (takes)
  - At a global start time, machine can start a new job if idle
- Goal: minimize total idle time

Significance

- Intended to provide a foundation for quantitatively evaluating adaptation in learning systems
- Potential to impact how ML algorithms are implemented and deployed

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

SAND 2016-xxxP