Exceptional service in the national interest

Sandia National Laboratories

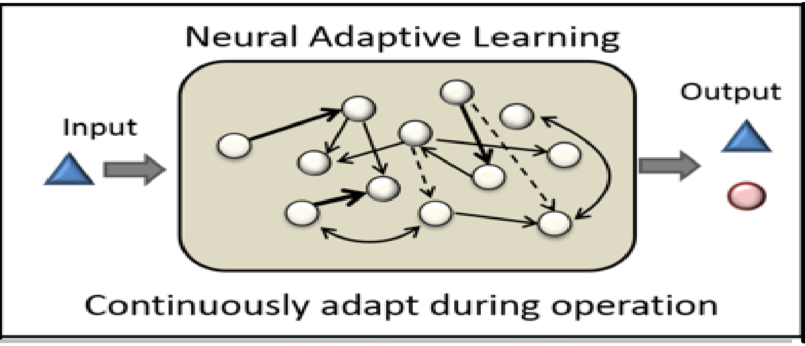

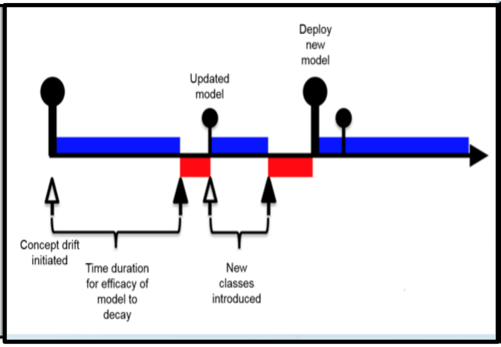

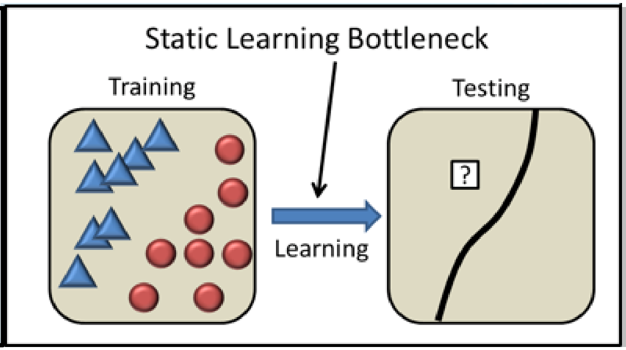

Studying Adaptive Learning through Game-Theoretic Modeling

Craig M. Vineyard, PhD

U.S. DEPARTMENT OF ENERGY
NNSA National Nuclear Security Administration
CCR Center for Computing Research

Sandia National Laboratories is a multi-mission 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 NO. 2011-XXXXP

Adaptive Learning

- The learning phase of an algorithm addresses the mechanism by which adjustments are made in the learning process (such as weight tuning in a neural network)

- One of the differentiating capabilities of the brain is continuous learning

- So the question becomes where are we with respect to machine learning?
  - Most data-driven algorithms in ML do not continuously adapt

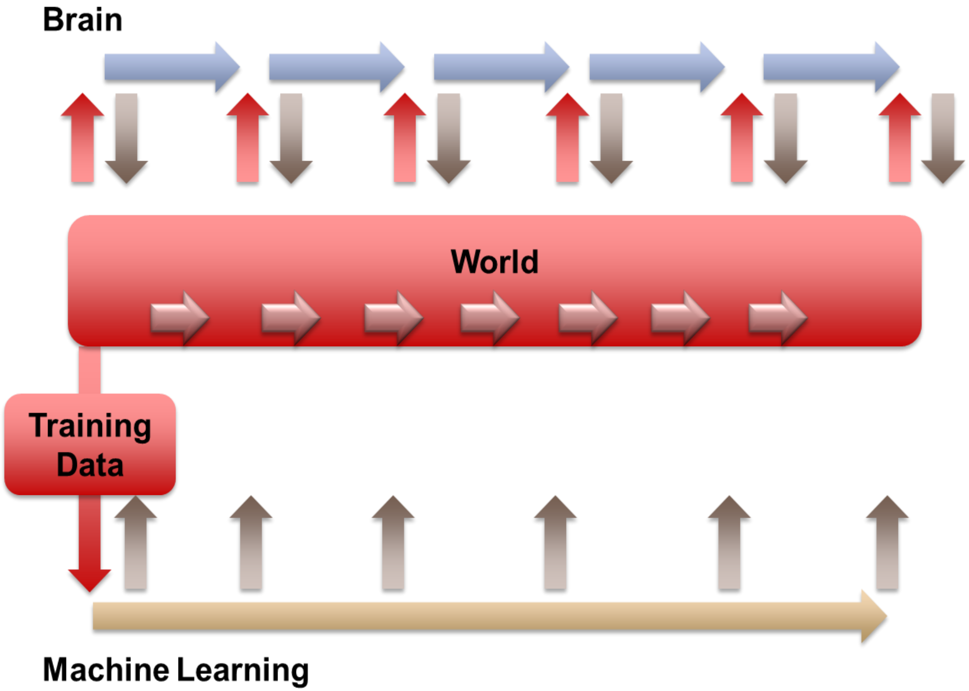

ML Learning Paradigms

Graph
Stacking/Super
Deep
Ensemble
Model
Top-down
Statistical
Learning
training
Unsupervised
Neural
Offline
Bayesian
Batch
Supervised
One-shot
Semi-supervised
Network
Inference
Bottom-up
learning
Dictionary
Online

…but they have limitations

Static Learning Bottleneck

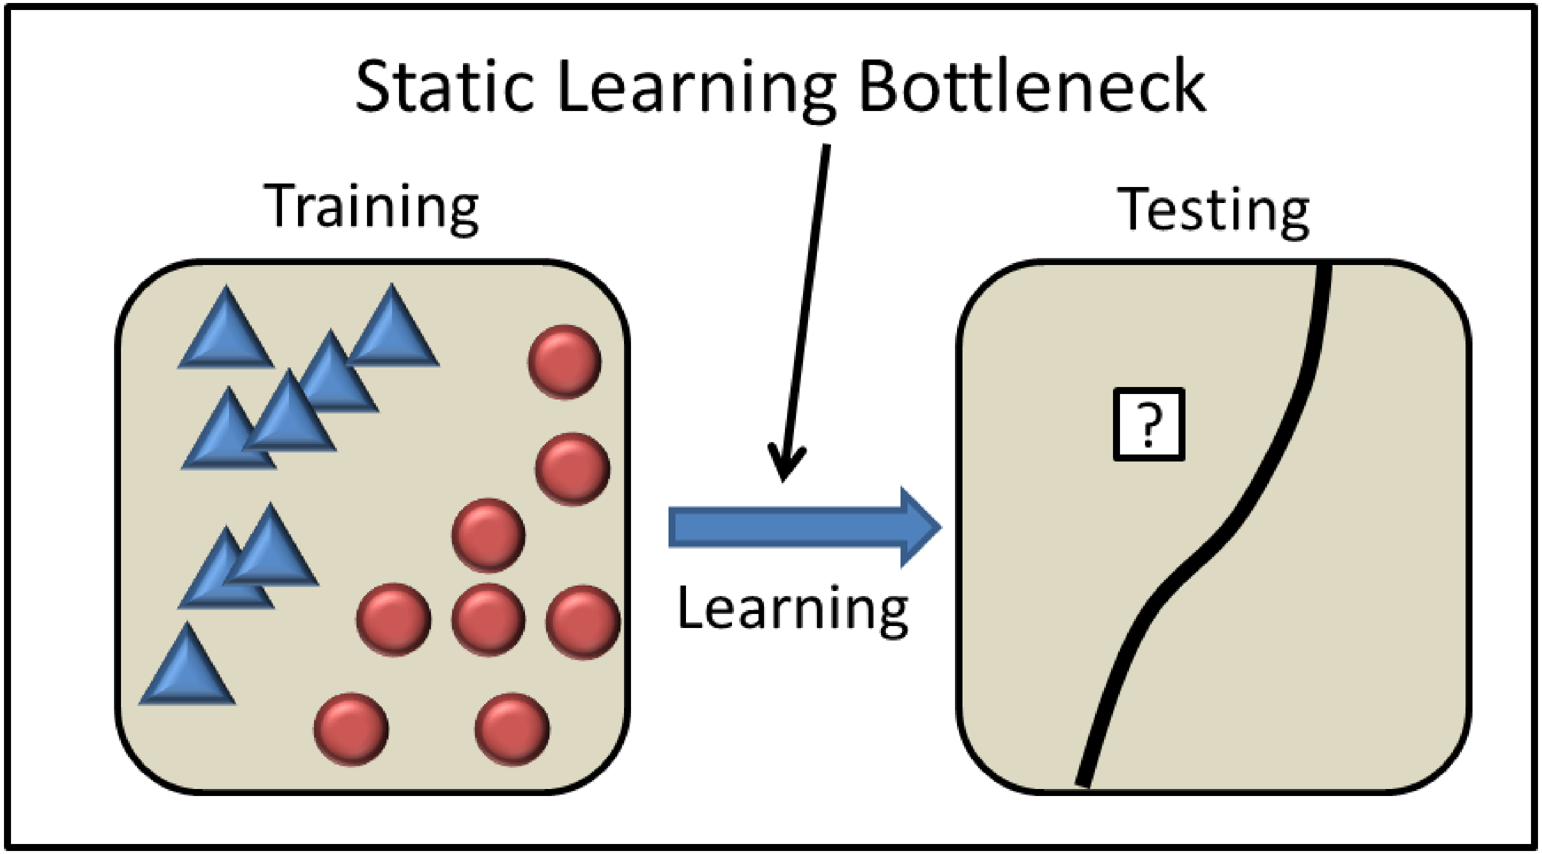

Continuous Neural Adaptation

- Synaptic plasticity
  - Dynamic alteration of the strength of the connections between neurons
- Structural plasticity
  - Addition and eliminations of neural network infrastructure

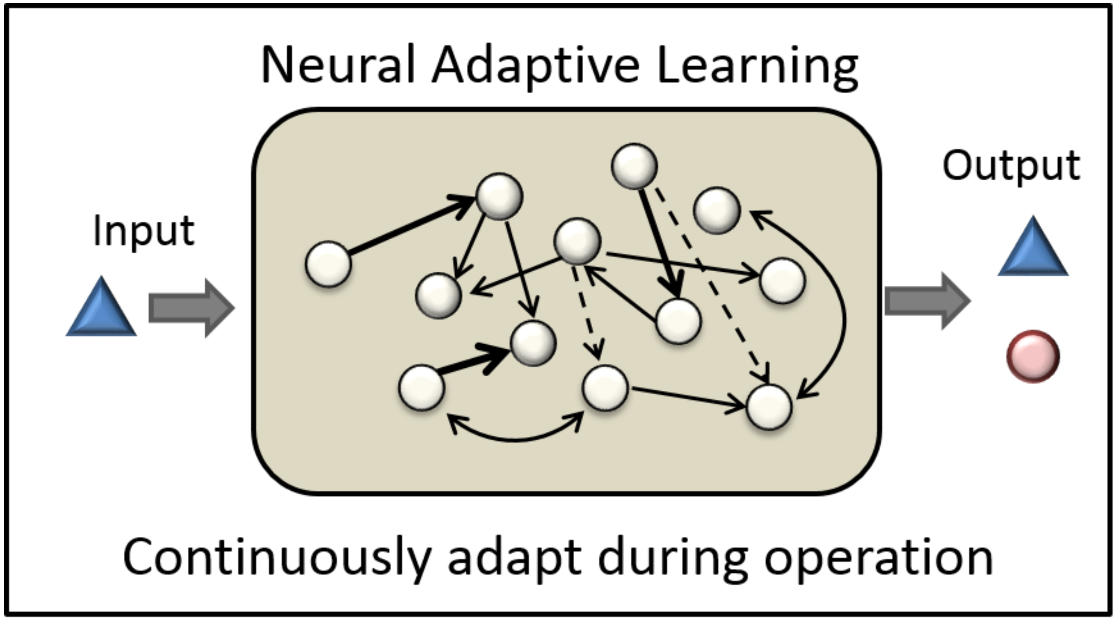

Game Theory

- Game theory is a branch of applied mathematics to formally analyze the strategic interaction between competing players

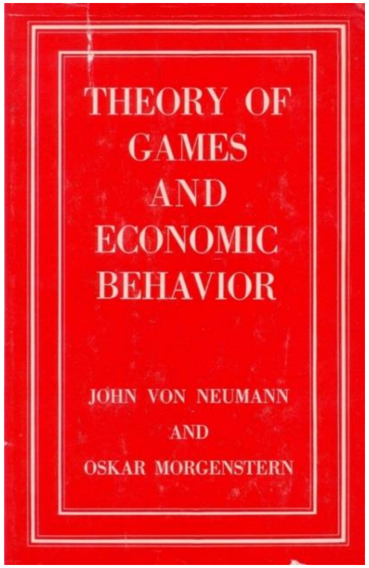

THEORY OF GAMES AND ECONOMIC BEHAVIOR

JOHN VON NEUMANN AND OSKAR MORGENSTERN

- Algorithmic Game Theory: the intersection of game theory & computer science
  - Analysis - analyzes algorithms from game-theoretic perspective, focus on properties such as equilibria
  - Design - focuses upon development of algorithms with desirable theoretical properties

Why game theory?

- Desirable properties for ML:
  - Leads to distributed computing, low overhead, simplicity, & provides a strategic perspective

Moving Target Defense (MTD)

- Use randomization, diversity, or change to make a computer system more difficult to attack (make it a “moving target”).
  - Randomized secret such as address-space layout randomization
  - Reset environment: new passwords, micro reboot, etc.
  - Deploy decoys. Change the real vs. decoys.

KEY:

- There is some information that helps the attacker as (s)he acquires it (e.g. in attempting to attack a system)
- The defender can take this information away, at least temporarily

PLADD

- Probabilistic Learning Attacker Dynamic Defender (PLADD)
  - Extension of FlipIt attacker and defender model
- Two players & one contested resource
- 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

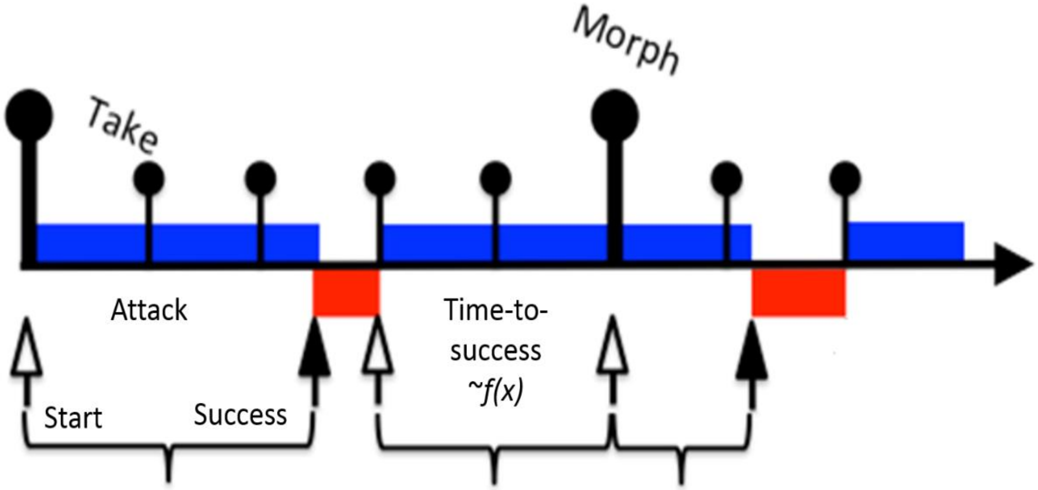

- Strategy: when to move?

PLADD for Learning = FLANEL

- Fundamental Learning Algorithm aNalysis and Exploration of Limits (FLANEL)
  - Modest extension that adds considerable complexity

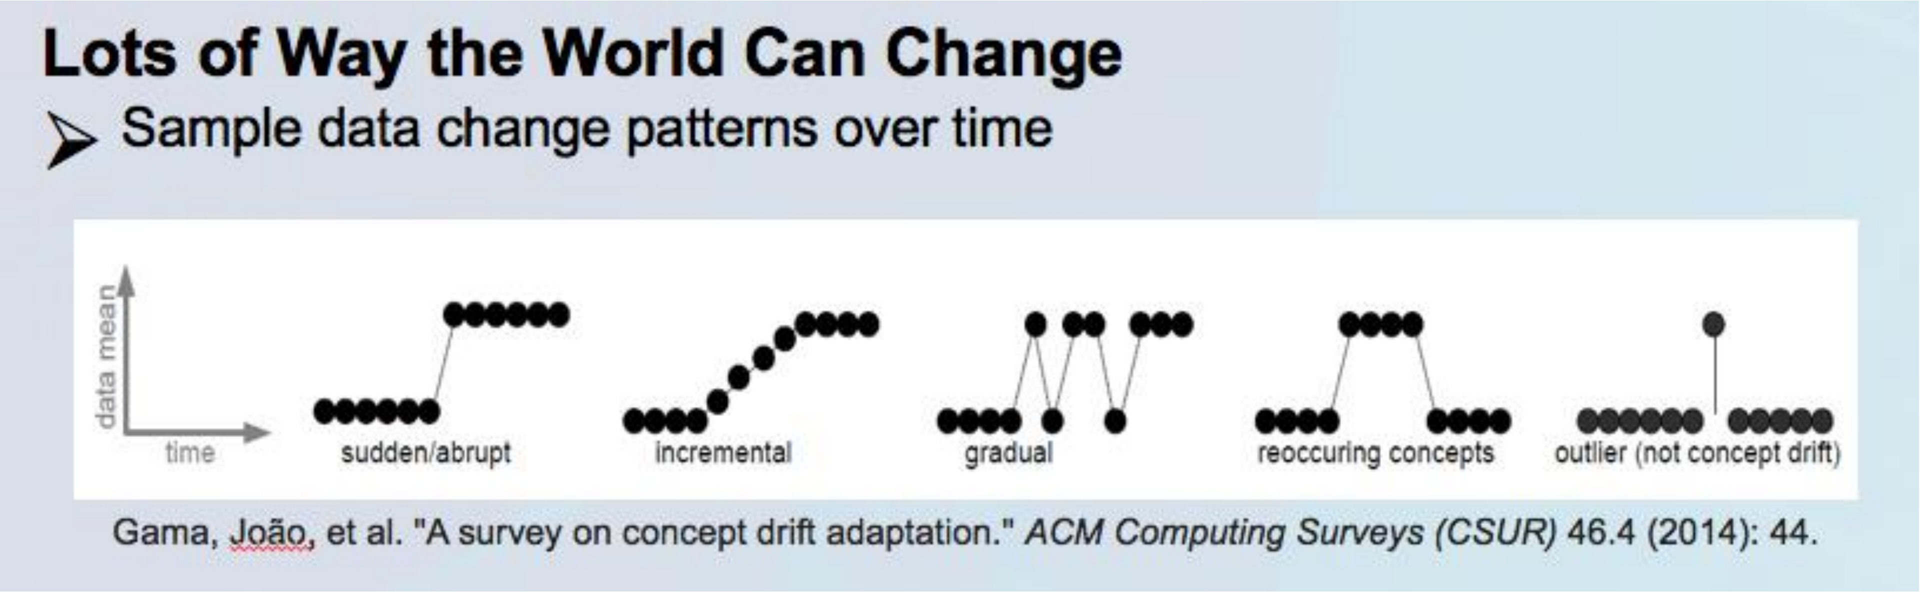

Lots of Way the World Can Change

- Sample data change patterns over time

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

FLANEL

- Morph = Rebuild the system (e.g. classifier)
- Take = Short-term improvement

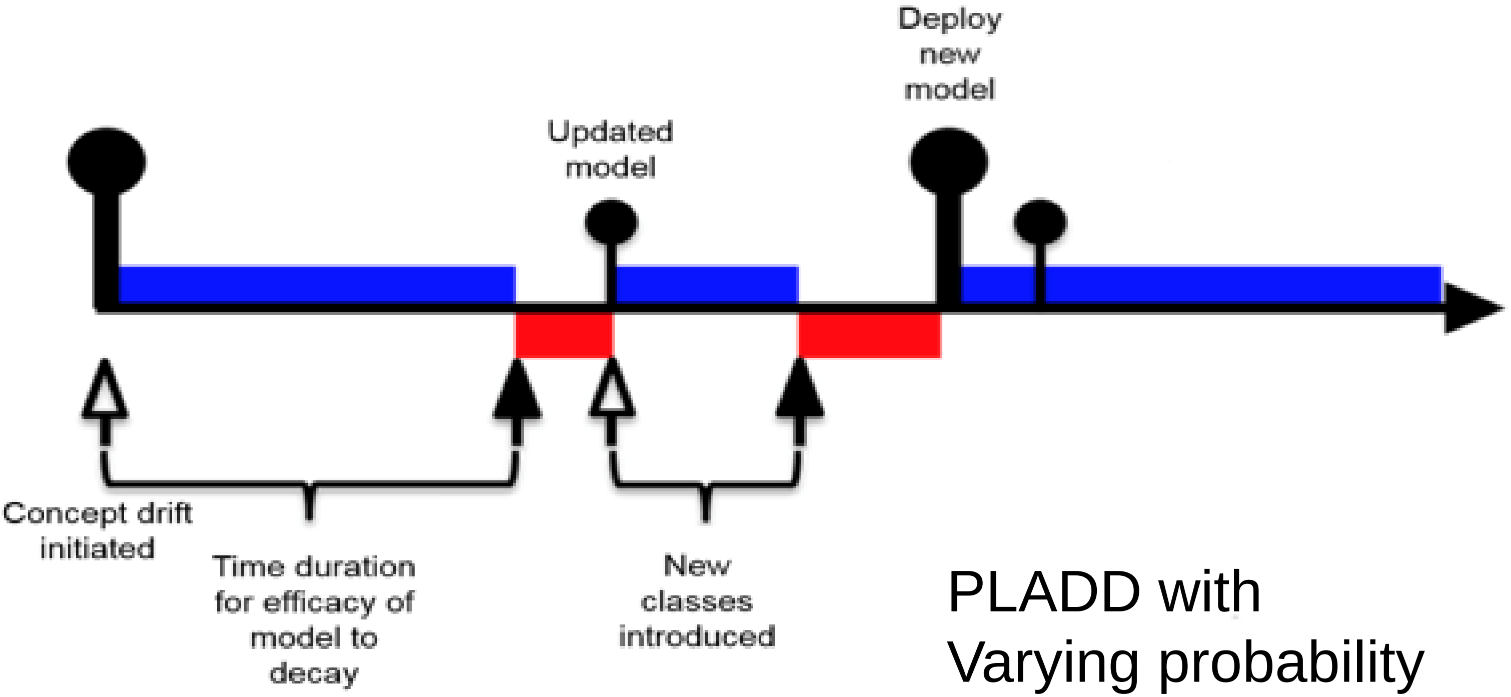

PLADD with
Varying probability
distributions

Exploring Alternatives to *Simulation –vs– Analytical*

- Analysis continuum

Simulation	Stochastic Programming	Analytical
Increasing Flexibility, Expressiveness		Increasing Generality

- Challenges:
  - Analytical: optimal response over continuous (infinite) parameters
    - May require restrictive / unrealistic assumptions (e.g., periodic moves)
  - Simulation: enumerate (subset of) parameters and collect statistics
    - Search by full enumeration frequently computationally intractable
- Opportunity:
  - Leverage numerical optimization to gain prescriptive insights while preserving much of the flexibility of simulation

Method 1: Stochastic Programming

- Key idea in stochastic programming:
  - approximate uncertainty by sampling outcomes
- Approximate attacker's strategy space by sampling possible random success-time outcomes
  - Attack scenarios
  - More scenarios gives a better approximation
- Optimize to determine the defender's single best strategy against ALL scenarios
  - Non-anticipative (only one solution for all attacks)
- Extensive form is a mixed-integer program (MIP)
- Can express more easily as a disjunctive program (DP)
  - Convert DP to MIP

Stochastic Programming Example

Idea:

- Study the time between two major model rebuilds (morphs)
- Fix the number of takes
- Draw many concrete instantiations

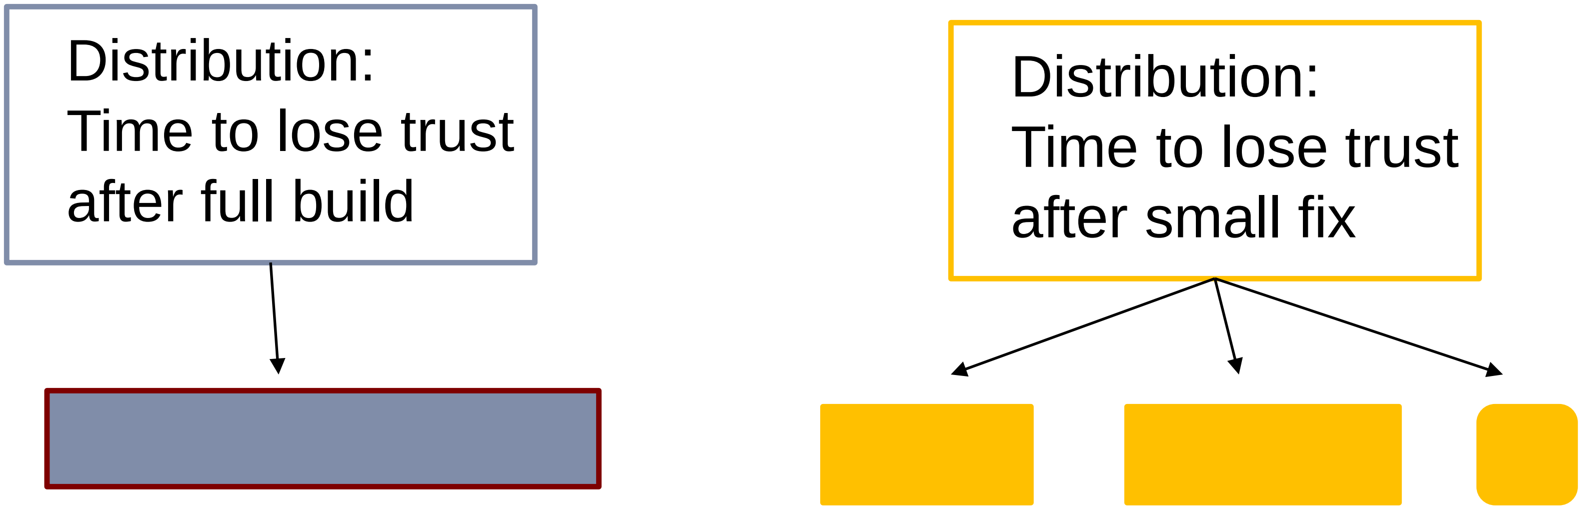

Distribution:
Time to lose trust
after full build

Distribution:
Time to lose trust
after small fix

Stochastic Program Example

- Given many concrete scenarios (Explicit time to model failure)

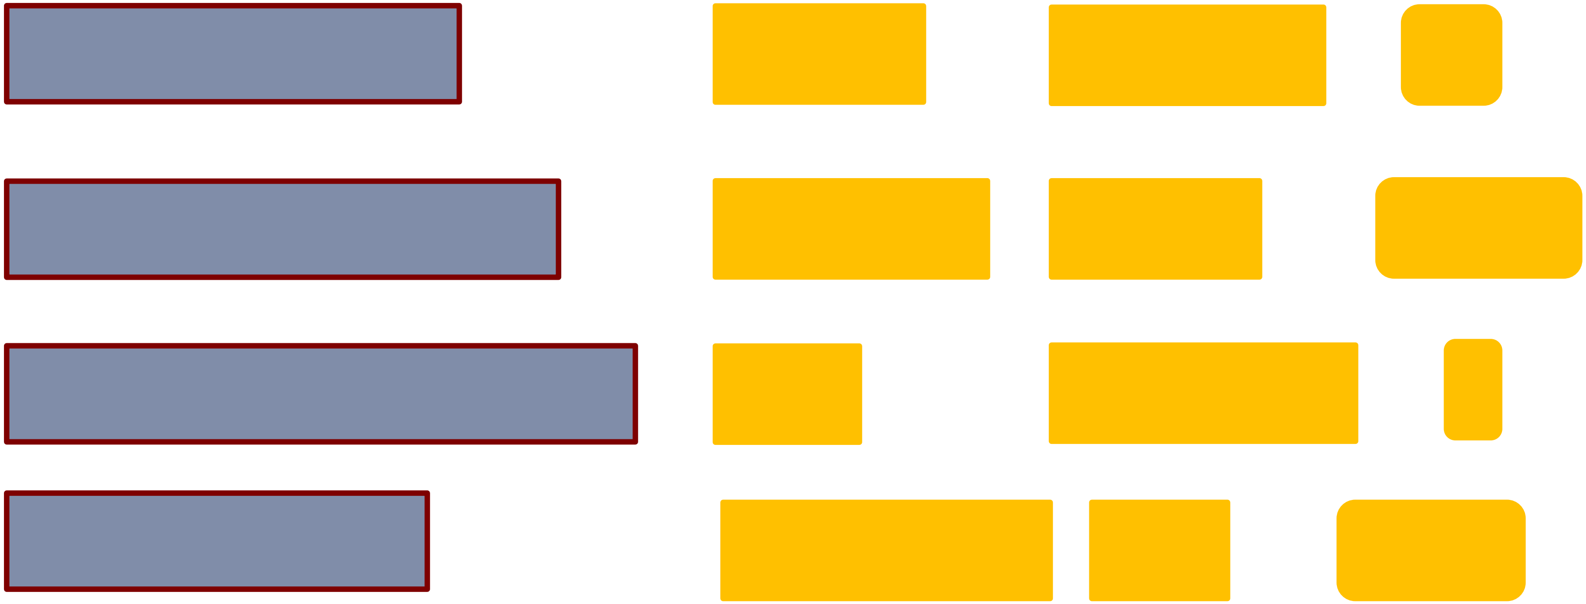

- Given only k (3 in this case) small fixes, when to do them?

Stochastic Program Example

- When to do the 3 small fixes?

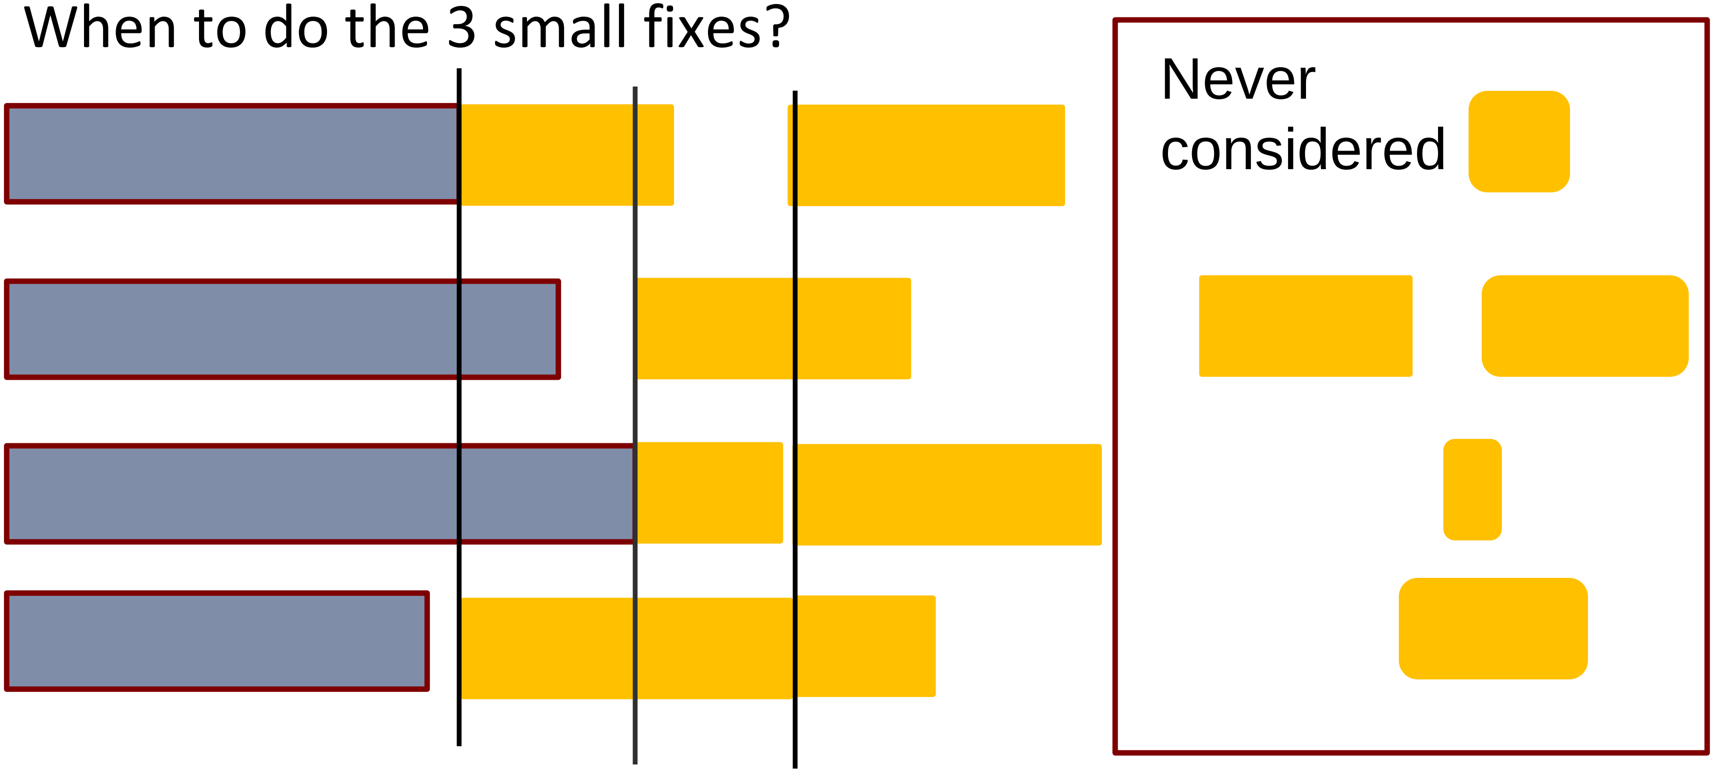

- Cost from the PLADD model: average time when you cannot trust the model.

FLANEL Cost

- When to do the 3 small fixes?

■ = model untrusted

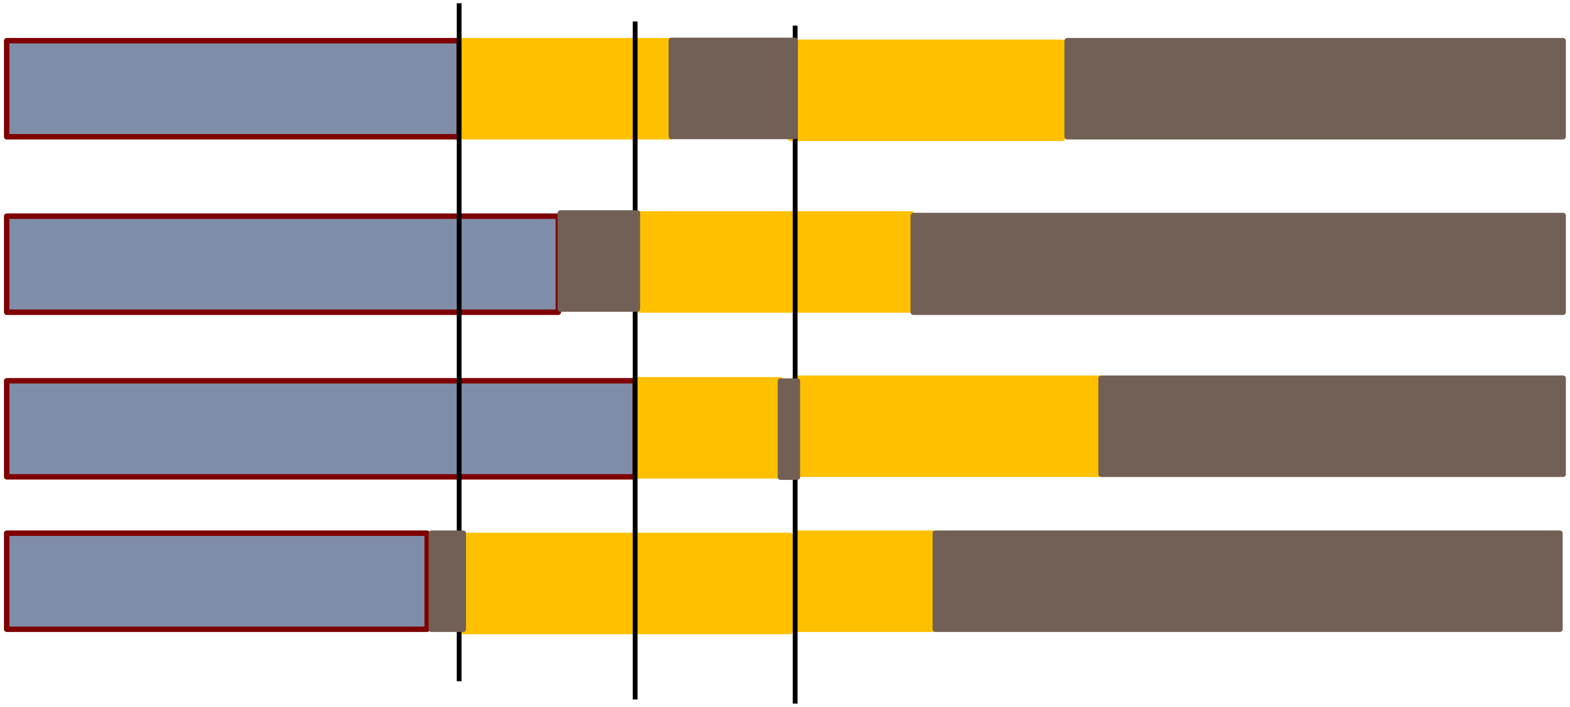

- Cost from the FLANEL model: average time when you cannot trust the model.

Method 2: Study Simpler Settings

- Streaming setting

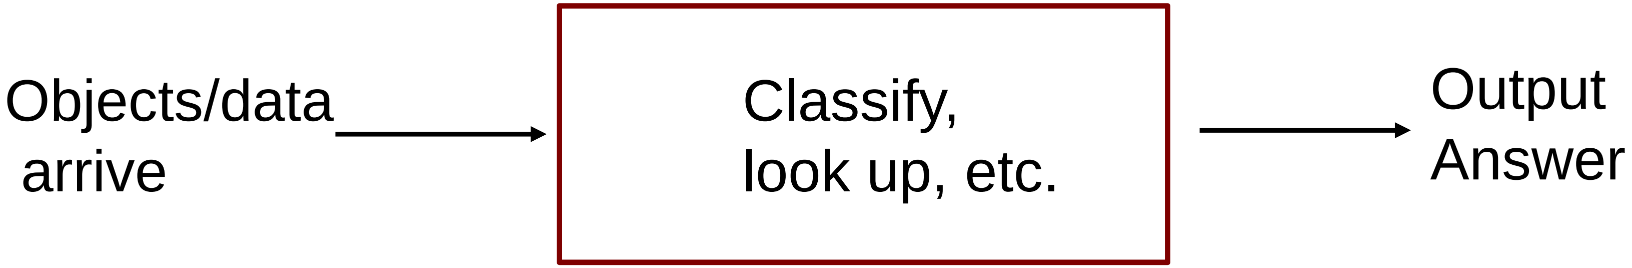

Objects/data arrive

Classify, look up, etc.

Output Answer

- Keep up with the stream
- When the data structures in the box are badly tuned, too slow
- Avoid dropping data elements

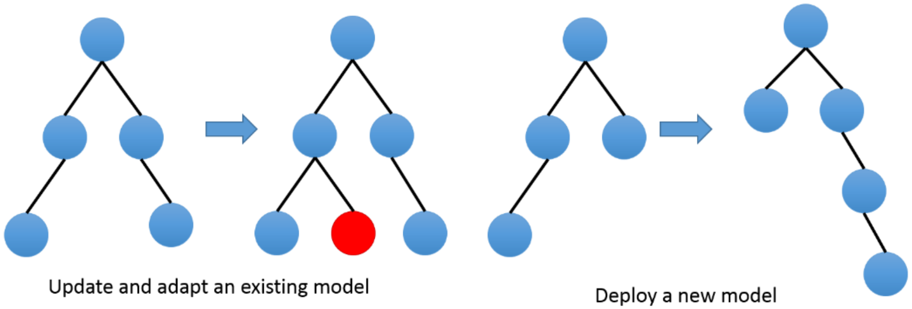

Update and adapt an existing model

Deploy a new model

Conclusion

- Static Learning Bottleneck – need for adaptive learning

- Working on a theoretical understanding of the problem
  - Need a holistic view not just Band-Aid solutions for individual problems
  - Mathematics of game theory are advantageous

- Presented FLANEL as an adaptive learning analysis framework
  - Intended to provide a foundation for quantitatively evaluating adaptation in learning systems
  - Potential to impact how ML algorithms are implemented and deployed

Thank you

Backup Slides

Interference

- Google’s DeepMind announced in February 2015 that they’d built a system that could beat 49 Atari games
  - However, each time it beat a game the system needed to be retrained to beat the next one

- "To get to artificial general intelligence we need something that can learn multiple tasks," says DeepMind researcher Hadsell. "But we can’t even learn multiple games."

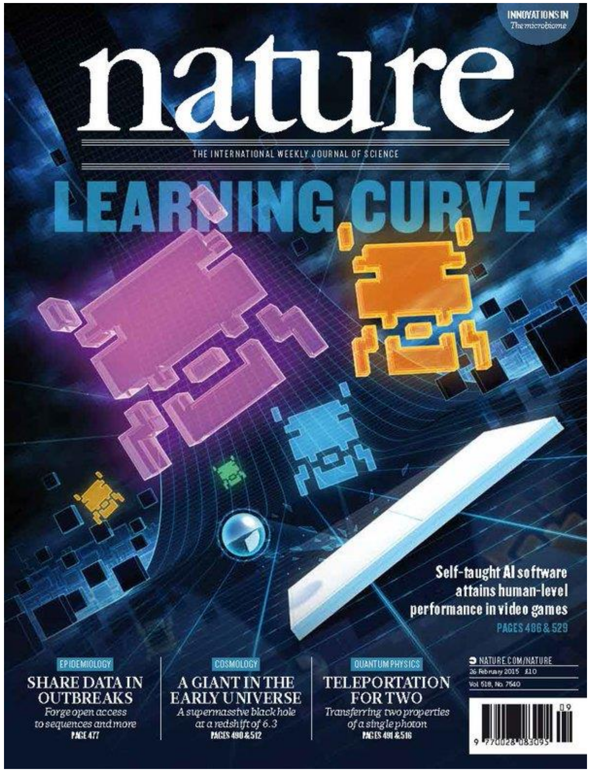
INNOVATIONS IN
The microbiome
nature
THE INTERNATIONAL WEEKLY JOURNAL OF SCIENCE
LEARNING CURVE
Self-taught AI software attains human-level performance in video games
PAGES 486 & 529
EPIDEMIOLOGY
SHARE DATA IN OUTBREAKS
Forge open access to sequences and more
PAGE 477
COSMOLOGY
A GIANT IN THE EARLY UNIVERSE
A supermassive black hole at a redshift of 6.3
PAGES 490 & 512
QUANTUM PHYSICS
TELEPORTATION FOR TWO
Transferring two properties of a single photon
PAGES 491 & 516
NATURE.COM/NATURE
26 February 2015 £10
Vol. 518, No. 7540

Nature Vol 518 Number 7540

Dynamic Environments

- Concept Drift: changes in the data over time
  - Virtual: changes in the underlying data distribution
  - Real: concepts themselves are changing

- Transfer Learning: the ability to utilize knowledge learned for one domain in learning a related but new domain

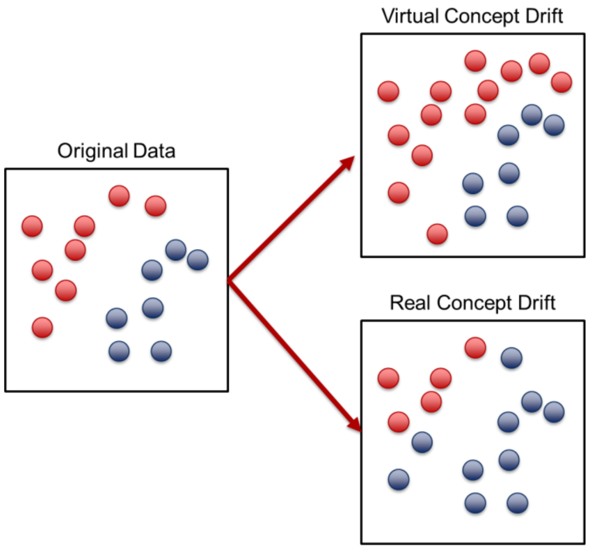

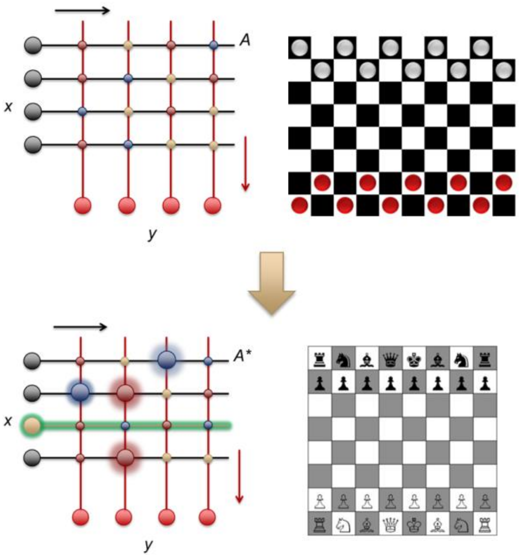

Key Limitation

> “The development of game theory in the early 1940s by John von Neumann was a reaction against the then dominant view that problems in economic theory can be formulated using standard methods from optimization theory. Indeed, most real world economic problems typically involve conflicting interactions among decision-making agents that cannot be adequately captured by a single (global) objective function, thereby requiring a different, more sophisticated treatment.”
>
> M. Pelillo and A. Torsello

- An analogous statement can be said about machine learning
  - Many learning problems involve dynamics that cannot be adequately capture by a single global objective function

FlipIt Example

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

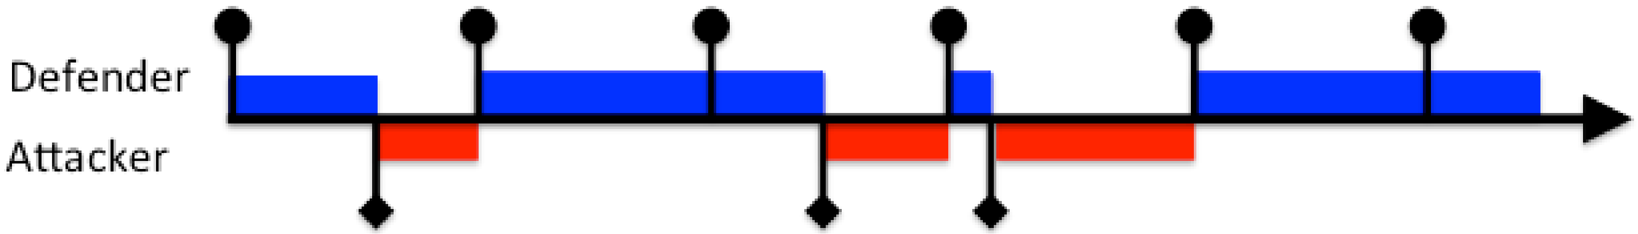

New Game: PLADD

- Probabilistic Learning Attacker Dynamic Defender
- Morphs reset to the start. Between morphs is a finite game
- With no morphs, the game is infinite, like FlipIt
- Difference between finite and infinite games is benefit of MTD

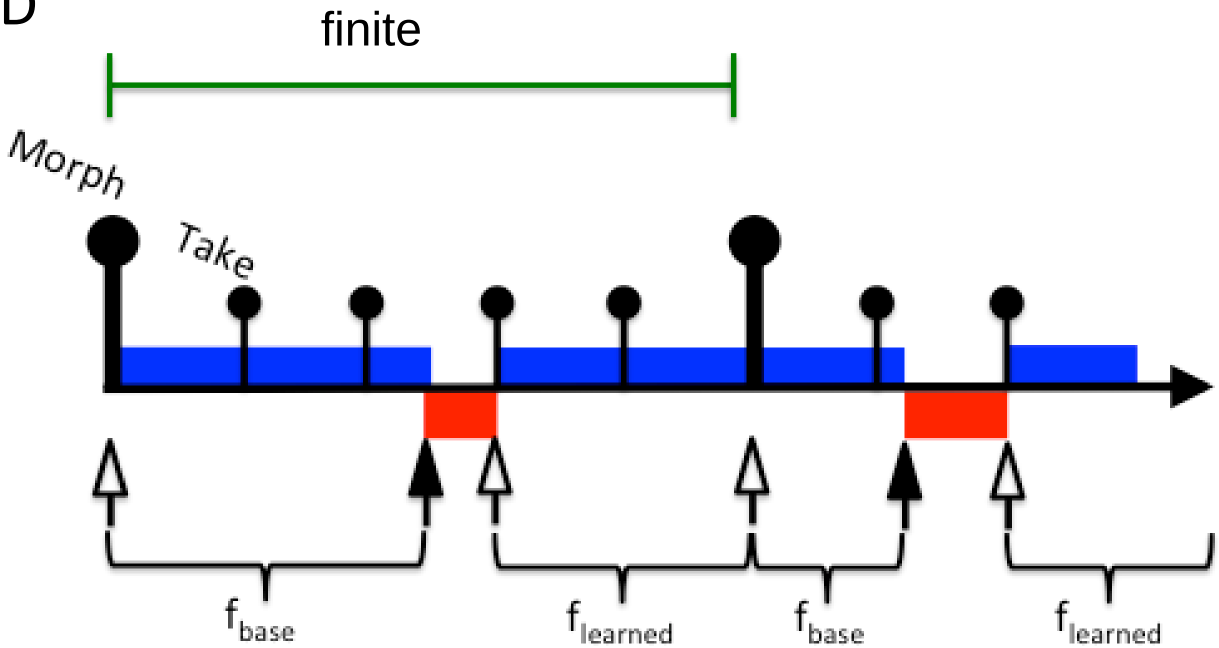

Formulating and Solving Stochastic Programs

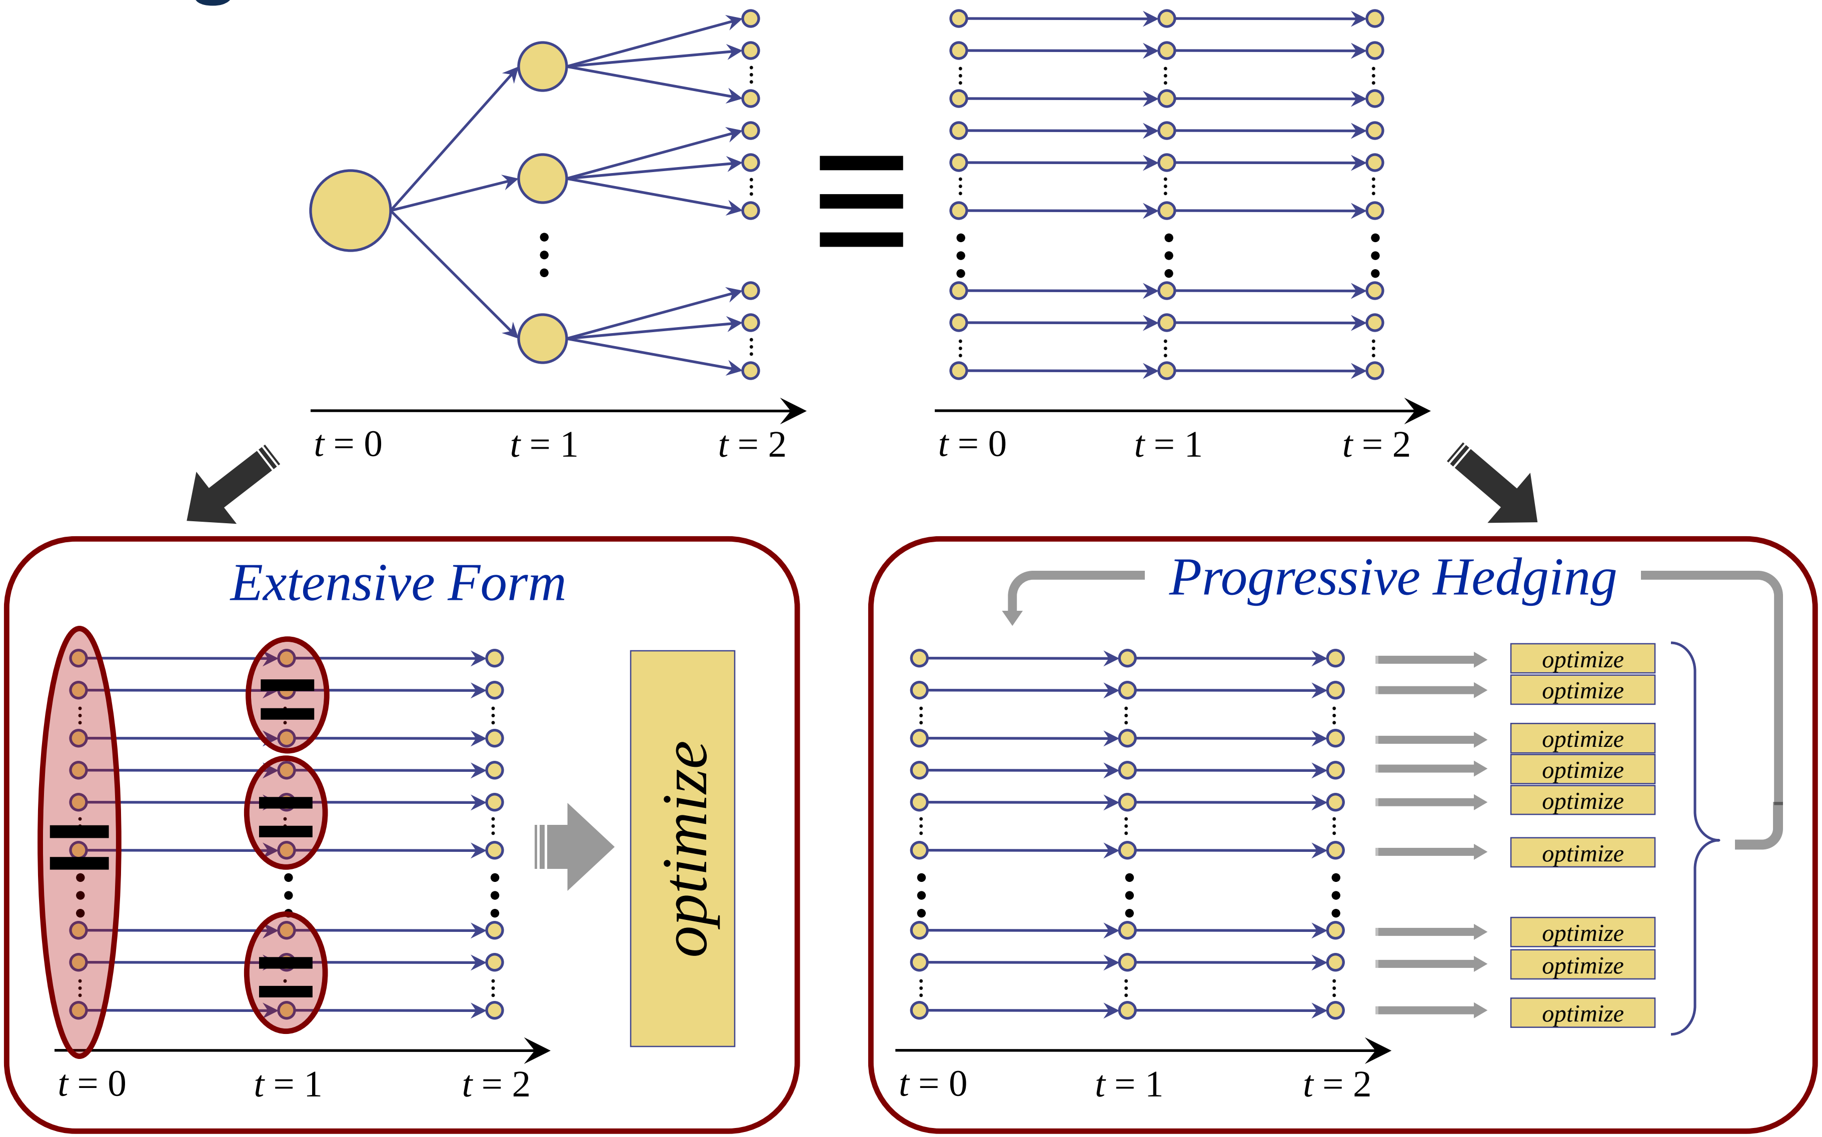

Extensive Form

Progressive Hedging

Extensive Formulation: MIP

- One schedule/strategy to minimize average cost

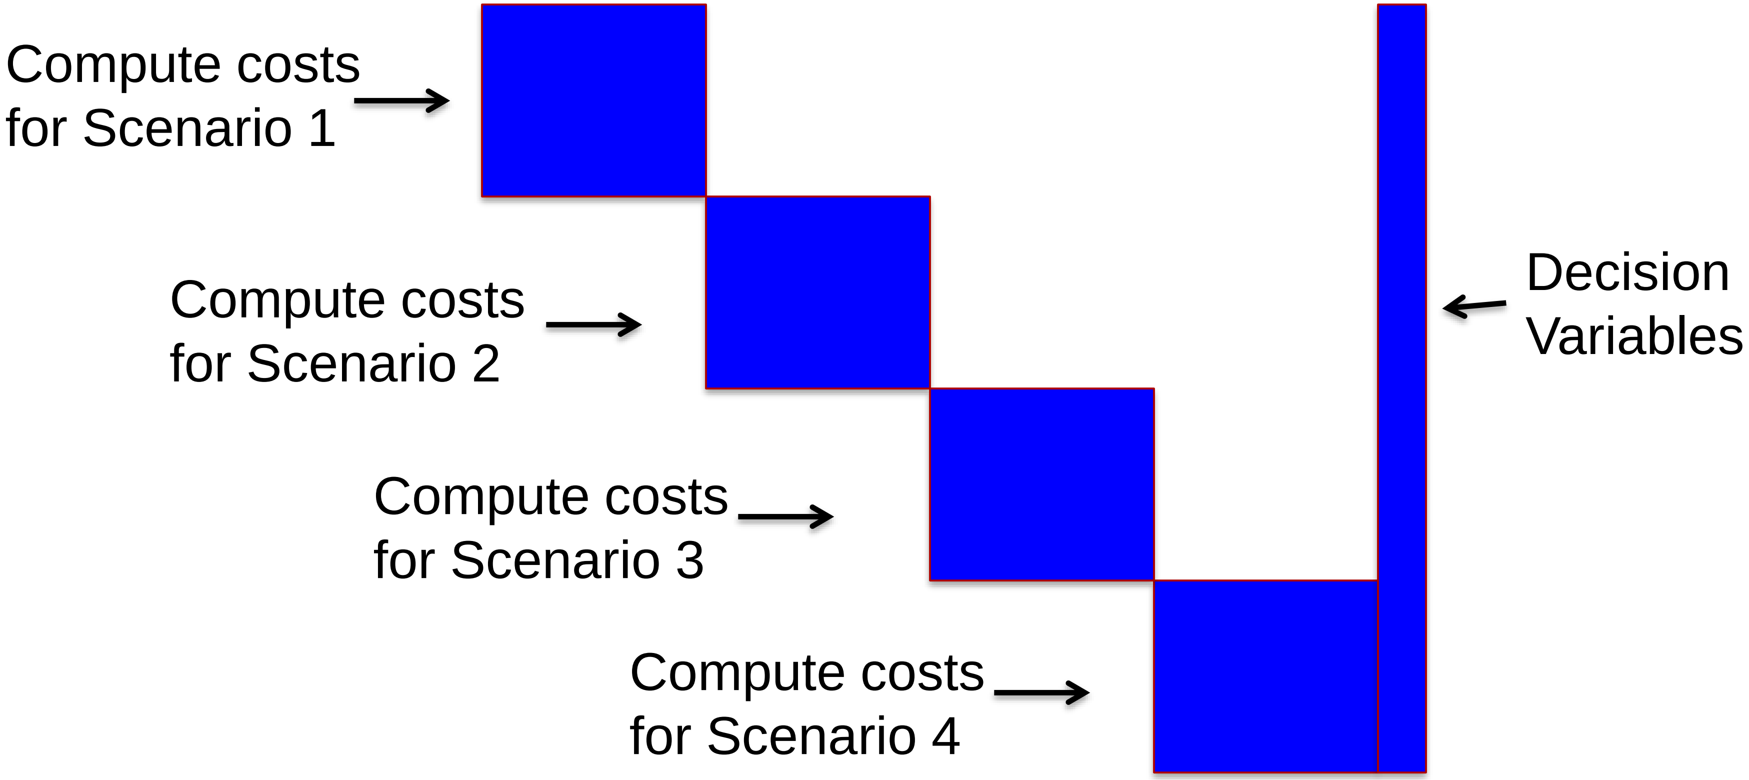
Compute costs for Scenario 1
Compute costs for Scenario 2
Compute costs for Scenario 3
Compute costs for Scenario 4
Decision Variables

Progressive Hedging

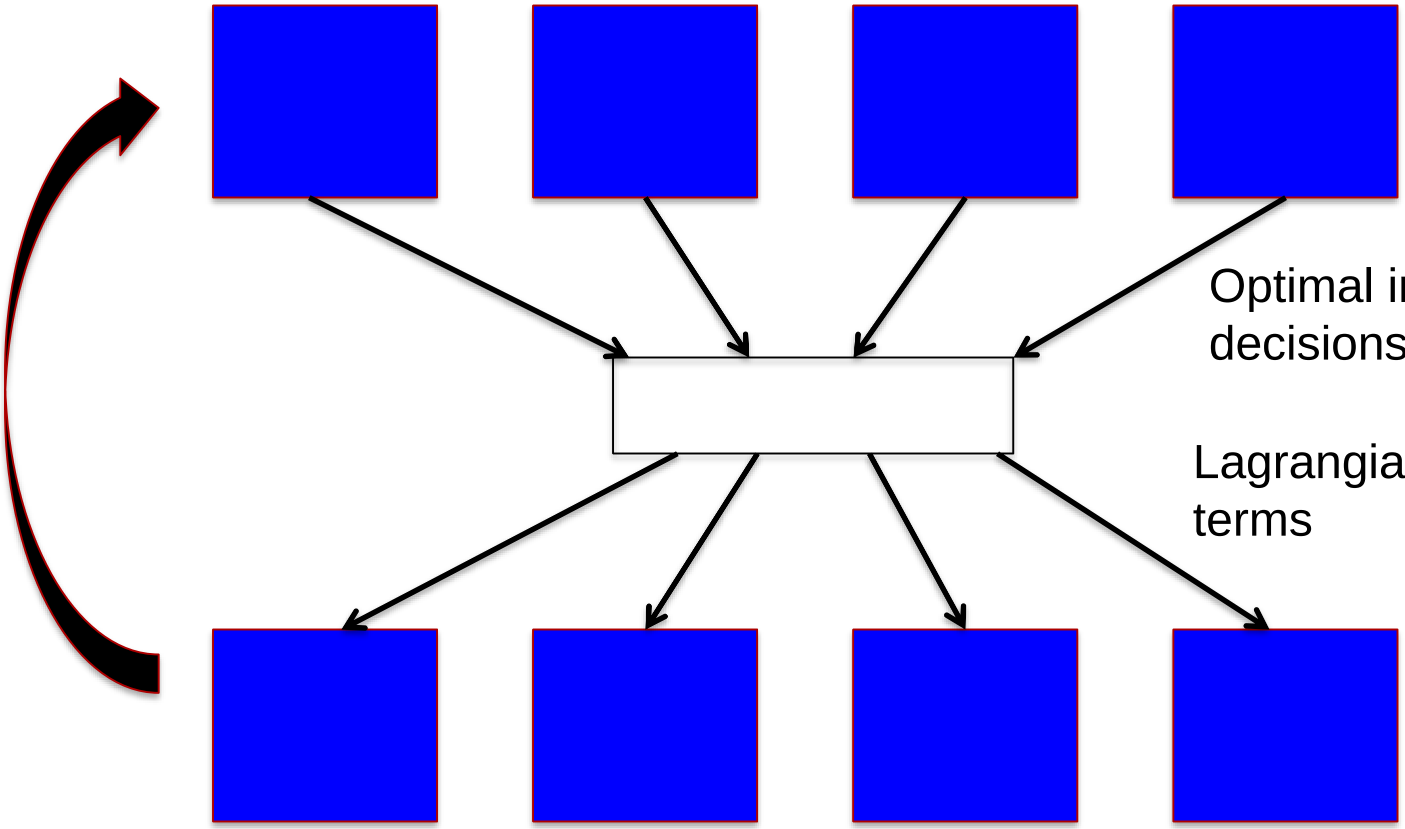

Optimal individual decisions

Lagrangian penalty terms