Real-time Decisions on Big Data

Leonard Newnham

NICE®

Overview

- Causata data platform
- Automated Decisioning objectives
- Challenges

What We Do and Where We Fit

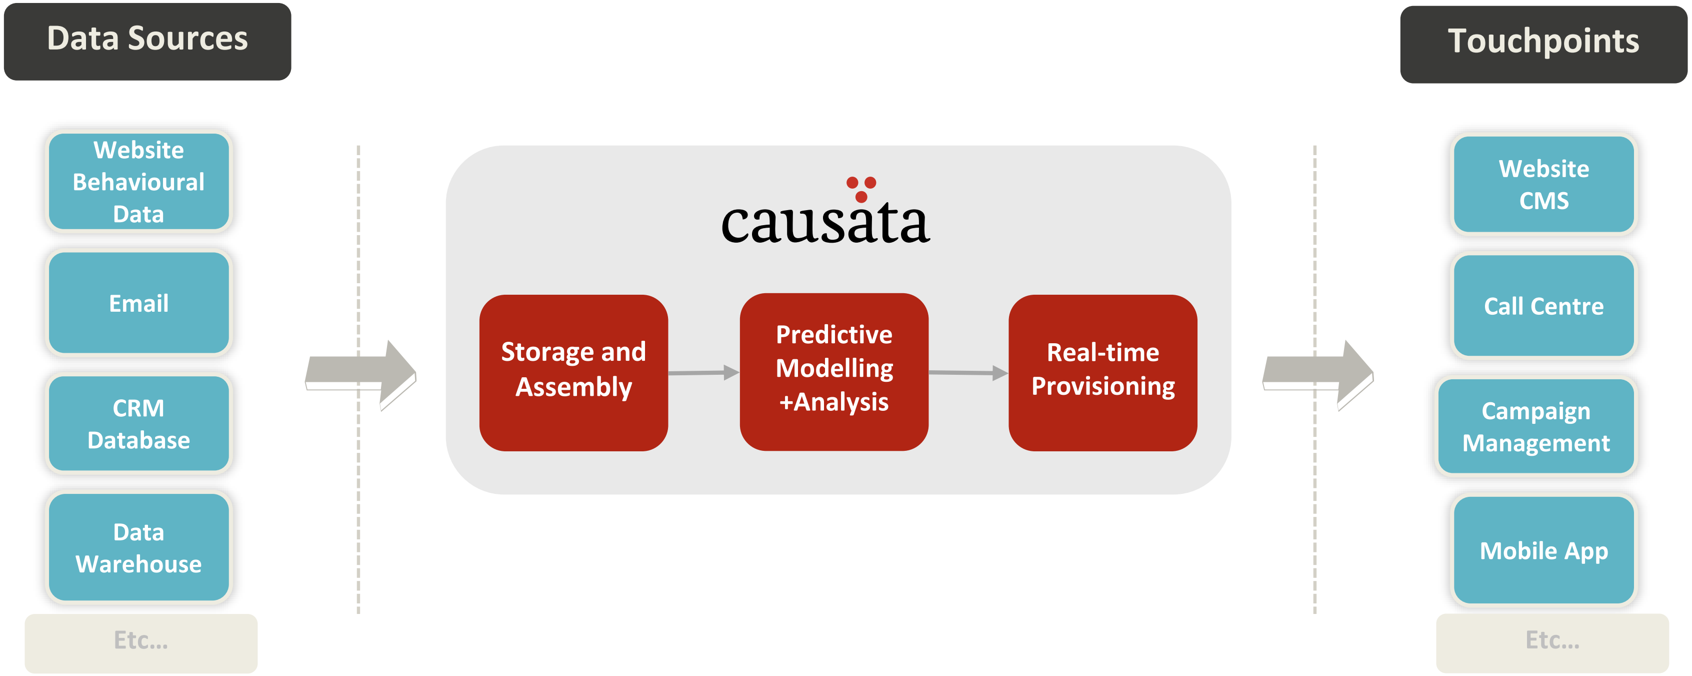

Customer Interactions

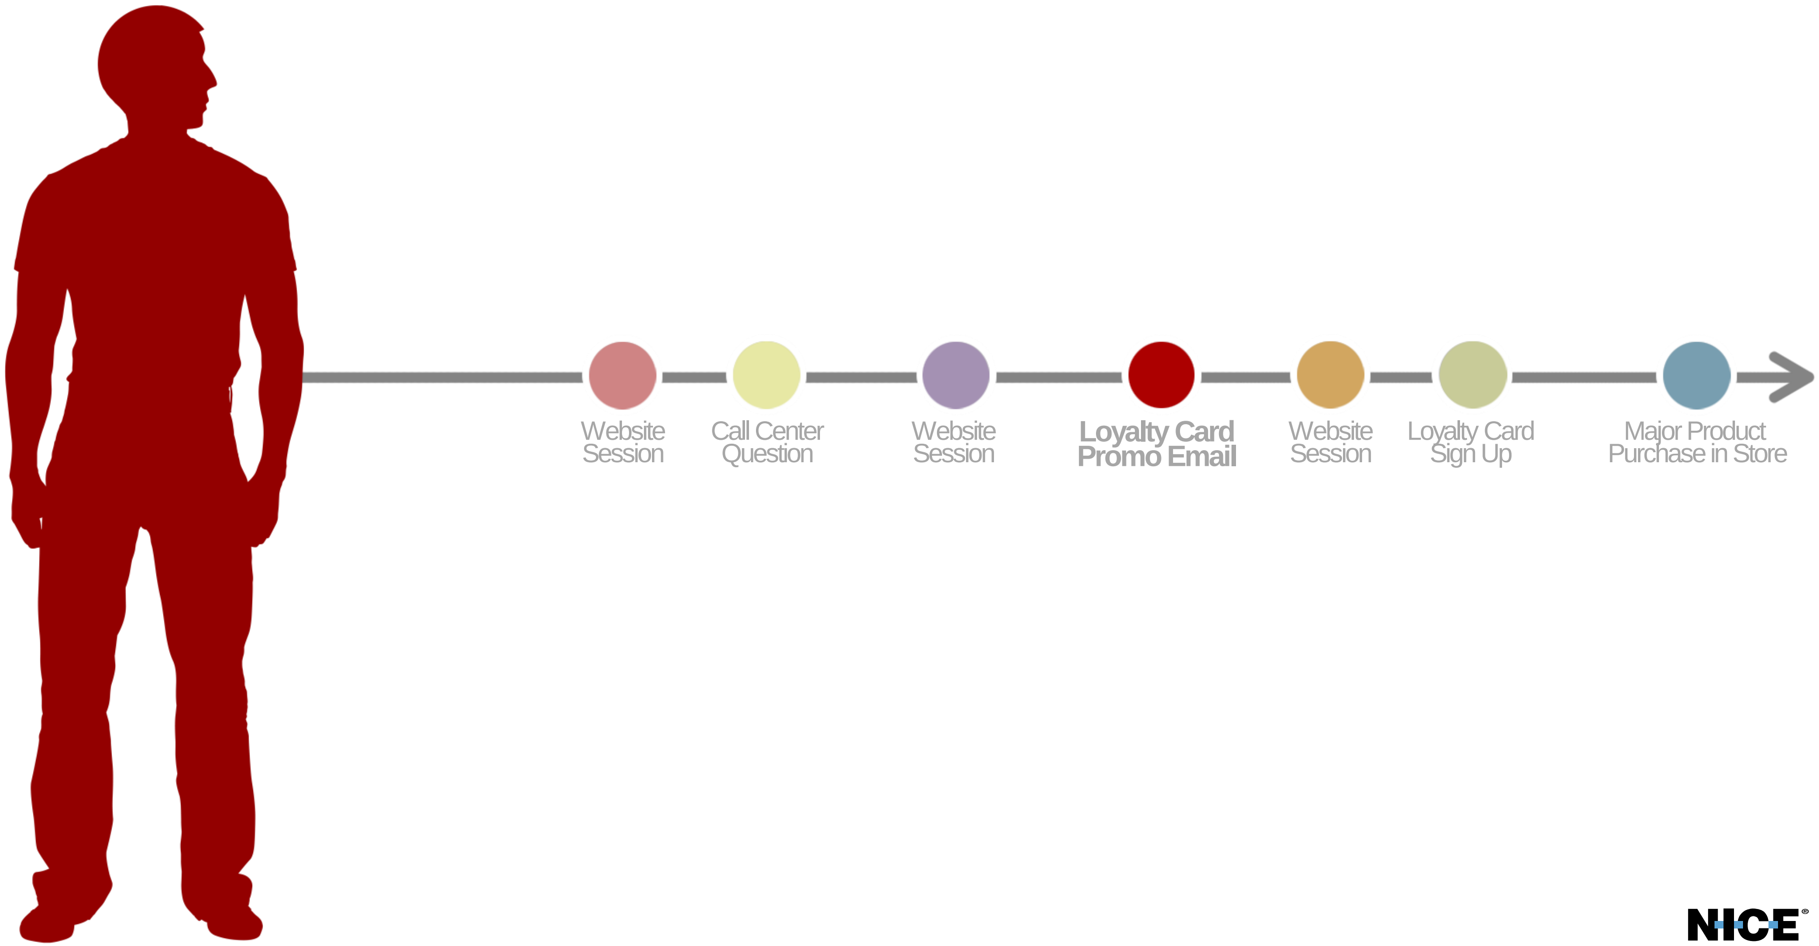

NICE

Variable Computation

Total spend in past month

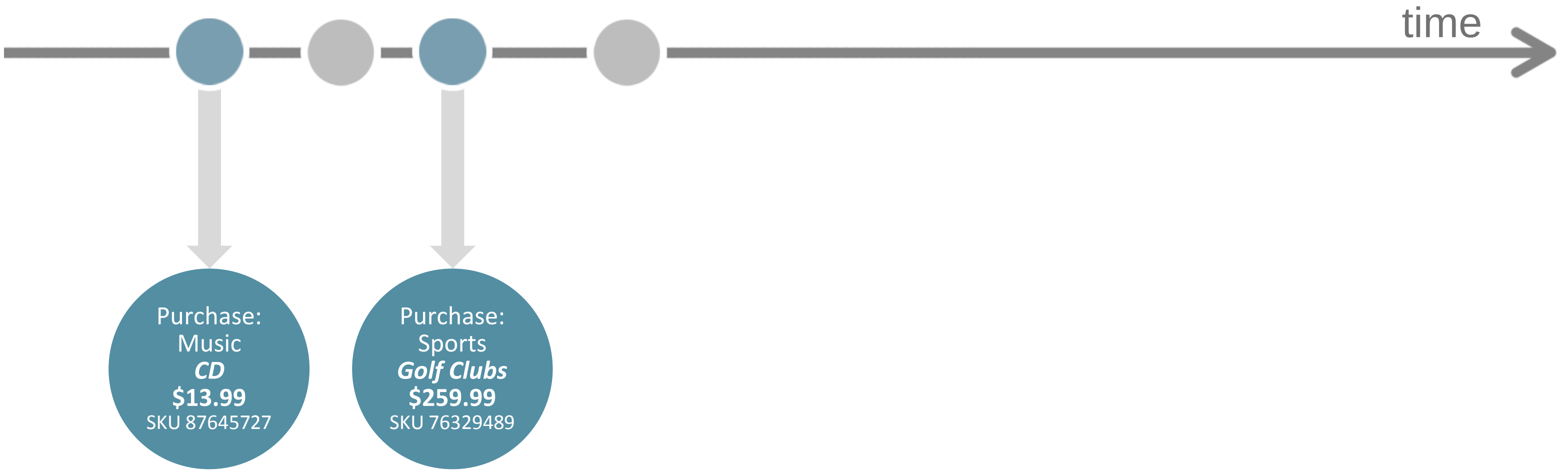

Select ***purchase events*** over past month and extract ***price***

Apply ***sum calculator***

Σ

13.99 + 259.99 = 273.98

Large Number of Variables

- Everything known about the visitor across multiple channels
  - Web data – page view history
  - Call centre data
  - Online accounts
  - Product holdings
  - geo-demographic data

Large Number of Variables

- Everything known about the visitor across multiple channels
  - Web data – page view history
  - Call centre data
  - Online accounts
  - Product holdings
  - geo-demographic data
- Meta-data on the available actions
  - price point, etc

Large Number of Variables

- Everything known about the visitor across multiple channels
  - Web data – page view history
  - Call centre data
  - Online accounts
  - Product holdings
  - geo-demographic data
- Meta-data on the available actions
  - price point, etc
- Environmental variables
  - Browser language, timezone, etc.

Large Number of Variables

- Everything known about the visitor across multiple channels
  - Web data – page view history
  - Call centre data
  - Online accounts
  - Product holdings
  - geo-demographic data
- Meta-data on the available actions
  - price point, etc
- Environmental variables
  - Browser language, timezone, etc.
- Typically 500+

Requirements

- Gathering data
  - Structure data by individual customer. Assemble all interactions, from many channels, into a complete customer record

Requirements

- Gathering data
  - Structure data by individual customer. Assemble all interactions, from many channels, into a complete customer record
- Making Decisions
  - Retrieve customer record and determine best action with low latency (< 100 ms)

Requirements

- Gathering data
  - Structure data by individual customer. Assemble all interactions, from many channels, into a complete customer record
- Making Decisions
  - Retrieve customer record and determine best action with low latency (< 100 ms)
- Displaying Results
  - Visualisation and explanation

Requirements

- Gathering data
  - Structure data by individual customer. Assemble all interactions, from many channels, into a complete customer record
- Making Decisions
  - Retrieve customer record and determine best action with low latency (< 100 ms)
- Displaying Results
  - Visualisation and explanation
- Scale to hundreds of millions of customers

Requirements

- Gathering data
  - Structure data by individual customer. Assemble all interactions, from many channels, into a complete customer record
- Making Decisions
  - Retrieve customer record and determine best action with low latency (< 100 ms)
- Displaying Results
  - Visualisation and explanation
- Scale to hundreds of millions of customers

- Optimise on goals that matter to customer

Optimise on Goals that Matter

Visit website from online banking → Research credit cards → Request application form → Speak to advisor in branch → Sign agreement → Ongoing relationship

Optimise on Goals that Matter

Visit website from online banking → Research credit cards → Request application form → Speak to advisor in branch → Sign agreement → Ongoing relationship

Current state of the art is to locally optimize each interaction, in terms of immediate next step

Optimise on Goals that Matter

Visit website from online banking → Research credit cards → Request application form → Speak to advisor in branch → Sign agreement → Ongoing relationship

Current state of the art is to locally optimize each interaction, in terms of immediate next step

With all the data, can optimize over true ***long term business goals***

Why Next Best Action May Not Always be Best

- To maximise long term gain there may be a short term cost:

Why Next Best Action May Not Always be Best

- To maximise long term gain there may be a short term cost:
  - special offers and discounts

Why Next Best Action May Not Always be Best

- To maximise long term gain there may be a short term cost:
  - special offers and discounts

- Or a deferred selling opportunity:

Why Next Best Action May Not Always be Best

- To maximise long term gain there may be a short term cost:
  - special offers and discounts

- Or a deferred selling opportunity:
  - restricted number of offers made after sale, eg
    - upgrades options after airline ticket purchase
    - extended warranty
    - Insurance
  - determine best time to send email

Reinforcement Learning

- Choose the actions which will yield the greatest long-term reward
- Reward can be any function we wish to optimize
- Rewards may be deferred to some time in the future

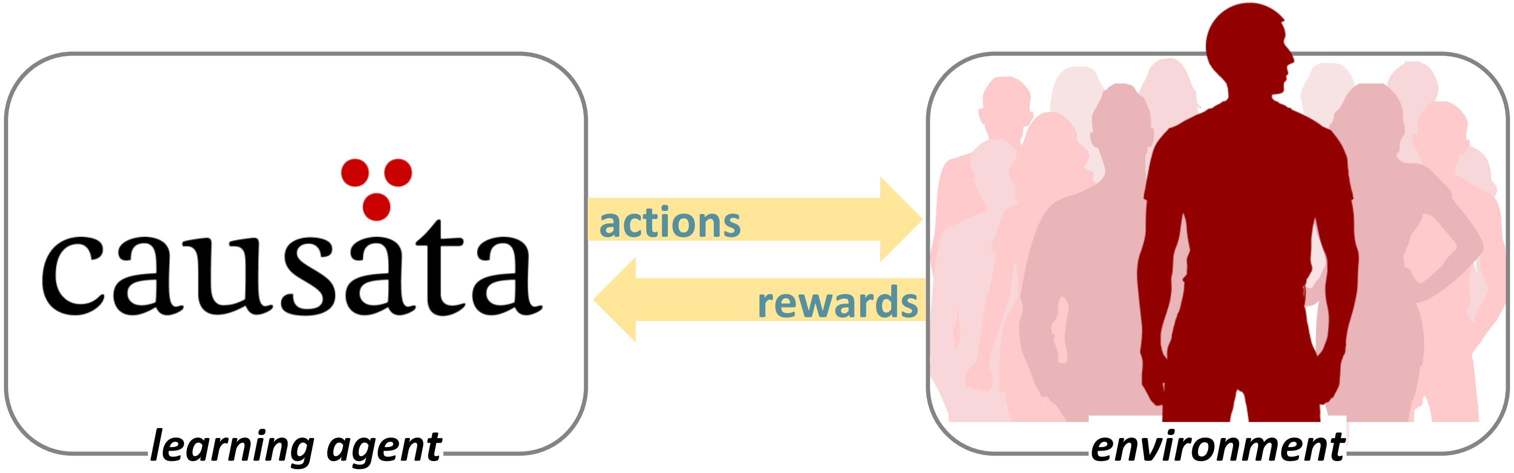

Challenges

- Big Data is not so big sometimes
- Concurrent customer interactions
- Speed of learning
- Visualisation
- Scalability

Big Data is not so Big Sometimes

- Learn objective functions that matter to client (part 2)

Big Data is not so Big Sometimes

- Learn objective functions that matter to client (part 2)
  - Revenue rather than click-through
  -

Big Data is not so Big Sometimes

- Learn objective functions that matter to client (part 2)
  - Revenue rather than click-through
  - Optimise multiple objectives

Big Data is not so Big Sometimes

- Changing Action Set Over Time

NICE

Big Data is not so Big Sometimes

- Changing Action Set Over Time
  - Individual actions may have a short lifetime
    - Marketing campaigns are frequently limited in time
    - May change due to seasonal variation – - eg summer sales
    - May change due to external factors – eg change of interest rate

Big Data is not so Big Sometimes

- Changing Action Set Over Time
  - Individual actions may have a short lifetime
    - Marketing campaigns are frequently limited in time
    - May change due to seasonal variation – - eg summer sales
    - May change due to external factors – eg change of interest rate

- -> We would like to not learn from scratch with every action change

Big Data is not so Big Sometimes

- Changing Visitor Behaviour Over Time

Big Data is not so Big Sometimes

- Changing Visitor Behaviour Over Time
  - Data becomes stale
  - It may be gradual
    - Popularity of product may wane
  - This may be abrupt
    - Summer heat-wave, Interest rate changes

Big Data is not so Big Sometimes

- Changing Visitor Behaviour Over Time
  - Data becomes stale
  - It may be gradual
    - Popularity of product may wane
  - This may be abrupt
    - Summer heat-wave, Interest rate changes

- At some point additional historic data will add more noise than signal

Big Data is not so Big Sometimes

- Changing Visitor Behaviour Over Time
  - Data becomes stale
  - It may be gradual
    - Popularity of product may wane
  - This may be abrupt
    - Summer heat-wave, Interest rate changes

- At some point additional historic data will add more noise than signal

- ->Learn from data up to that point and no more

Big Data is not so Big Sometimes

- Learn objective functions that matter to client
  - Revenue rather than click-through
- Changing action set over time
- Changing visitor behaviour over time

Big Data is not so Big Sometimes

- Learn objective functions that matter to client
  - Revenue rather than click-through
- Changing action set over time
- Changing visitor behaviour over time
- Learn rapidly in Proof of Concept implementations

Big Data is not so Big Sometimes

- Learn objective functions that matter to client
  - Revenue rather than click-through
- Changing action set over time
- Changing visitor behaviour over time
- Learn rapidly in Proof of Concept implementations
- Most data is noise

Big Data is not so Big Sometimes

- Learn objective functions that matter to client
  - Revenue rather than click-through
- Changing action set over time
- Changing visitor behaviour over time
- Learn rapidly in Proof of Concept implementations
- Most data is noise
- -> smart counting is not sufficient

Concurrent Customer Interactions

- At any time there are a large number of customer trajectories in progress

Concurrent Customer Interactions

- At any time there are a large number of customer trajectories in progress
- Time between each interaction is highly variable

Concurrent Customer Interactions

- At any time there are a large number of customer trajectories in progress
- Time between each interaction is highly variable

- -> Need to avoid long delays to learning

Concurrent Customer Interactions

- Need mechanism where:

Concurrent Customer Interactions

- Need mechanism where:
  - Learning can be transferred as quickly as possible *to other concurrent customers*

Concurrent Customer Interactions

- Need mechanism where:
  - Learning can be transferred as quickly as possible *to other concurrent customers*
  - Without waiting for next interaction or end of sequence

Speed of Learning

- Experience replay
- Regularisation
- Adaptive learning rate, eg IDBD
- Weight initialisation
- Improved Exploration
  - E-greedy
    - Simple but can perform badly when more exploration is required
  - UCB type exploration

Visualisation and Explanation

- How can I trust the system will work?
- What has the system learned?

Visualisation and Explanation

- How can I trust the system will work?
- What has the system learned?

- Simple accessible views
  - Not easy with complex internal representation

Visualisation and Explanation

- How can I trust the system will work?
- What has the system learned?

- Simple accessible views
  - Not easy with complex internal representation
- Explanation of why individual decisions are made

Visualisation and Explanation

- How can I trust the system will work?
- What has the system learned?

- Simple accessible views
  - Not easy with complex internal representation
- Explanation of why individual decisions are made

- Encourage operator engagement

Scalability and Redundancy

- Multiple learning agents
-

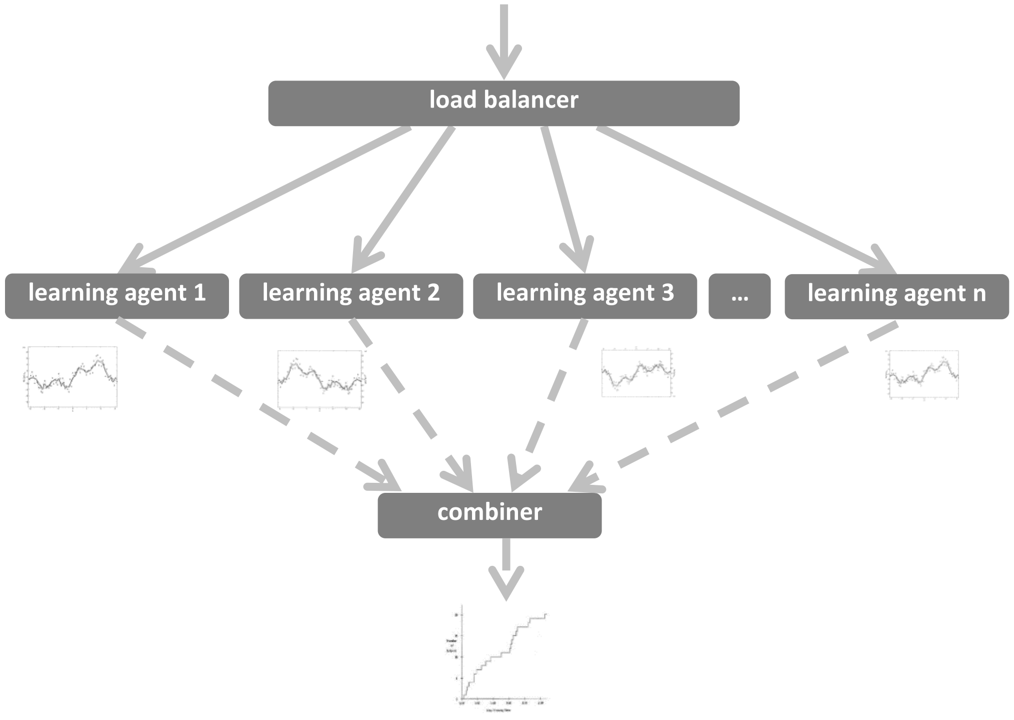

Scalability and Redundancy

- Multiple learning agents
- -> Regular dissemination of learning to other agents

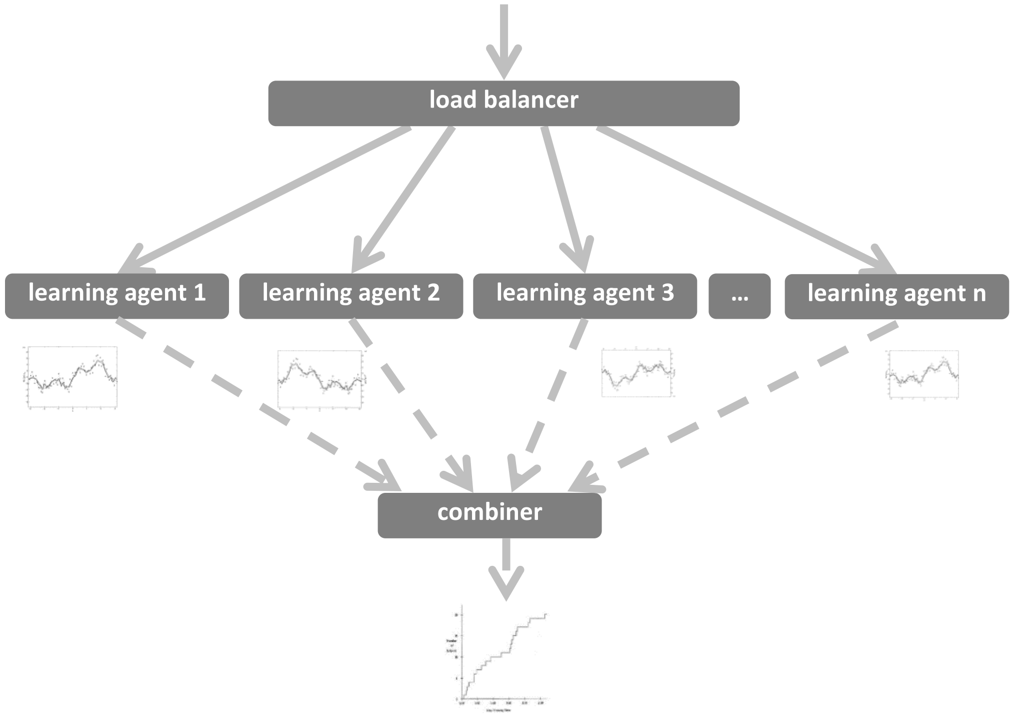

THANK YOU

Leonard Newnham
leonard.newnham@nice.com