Decision Camp 2017
Birbeck, University of London

Decision models for the Digital Economy

Vijay Bandekar
InteliOps Inc.

Agenda

- Problem Statement
- Proposed Solution
- Case studies and results
- Key takeaways
- Discussions

Digital Economy and Data Explosion

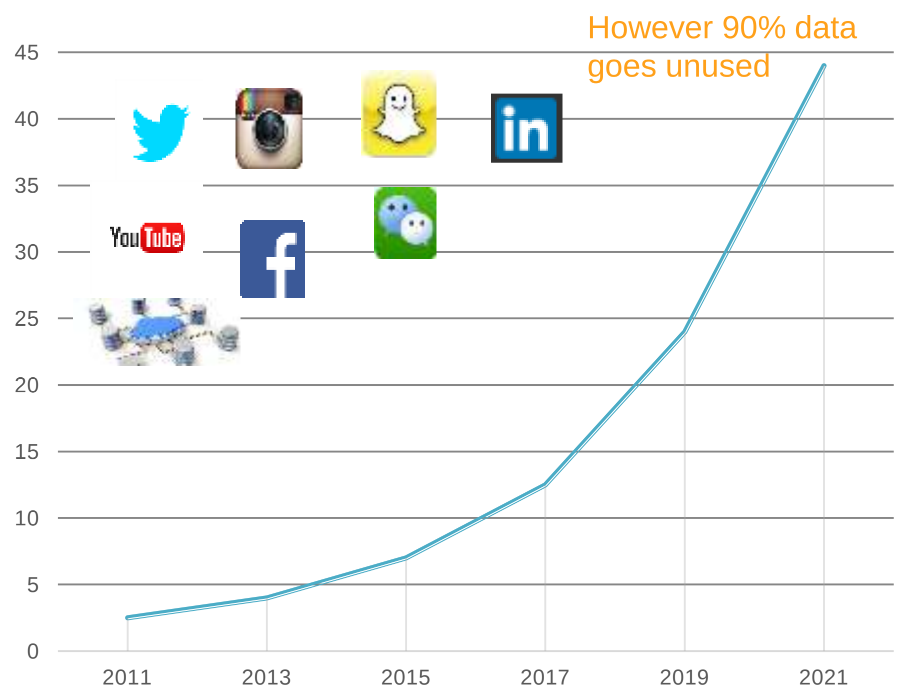

40 Zb of data by 2020, 44X in 11 Years
26 Bn IoT Devices by 2020

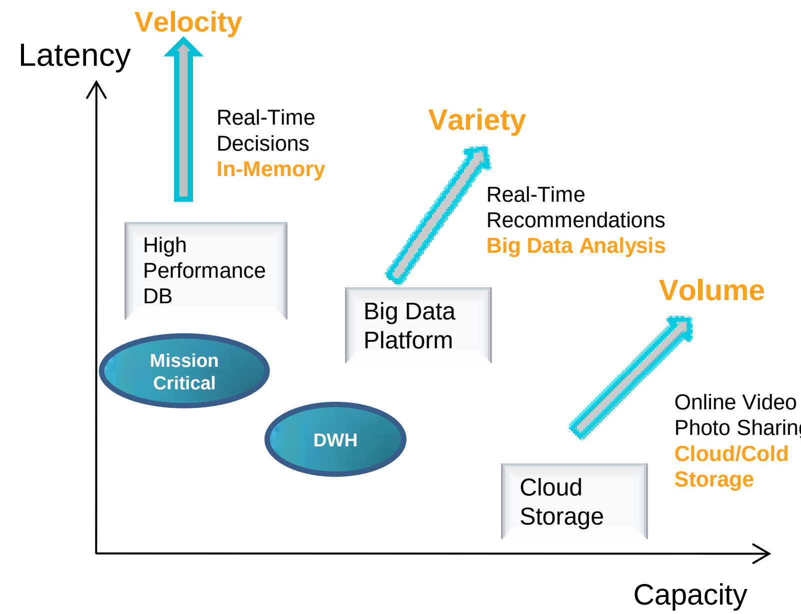

- Digital economy demands autonomous, self-learning, real-time decision making systems that ***sense***, ***comprehend*** and ***act***.
- Decisions must adapt to changing business environments.

Decision Models:
Inference, Reasoning and Deductions

What are decision models?	Business Logic Templates.
How do they work?	Workflows, rules engine, …
Available technologies?	Rete, logic programs, SQL

- Rules development and maintenance is cumbersome.
- Machine learning can generate rules/models, however, they must be manually incorporated into a business rules engines (BRE).

Reasoning and Inference

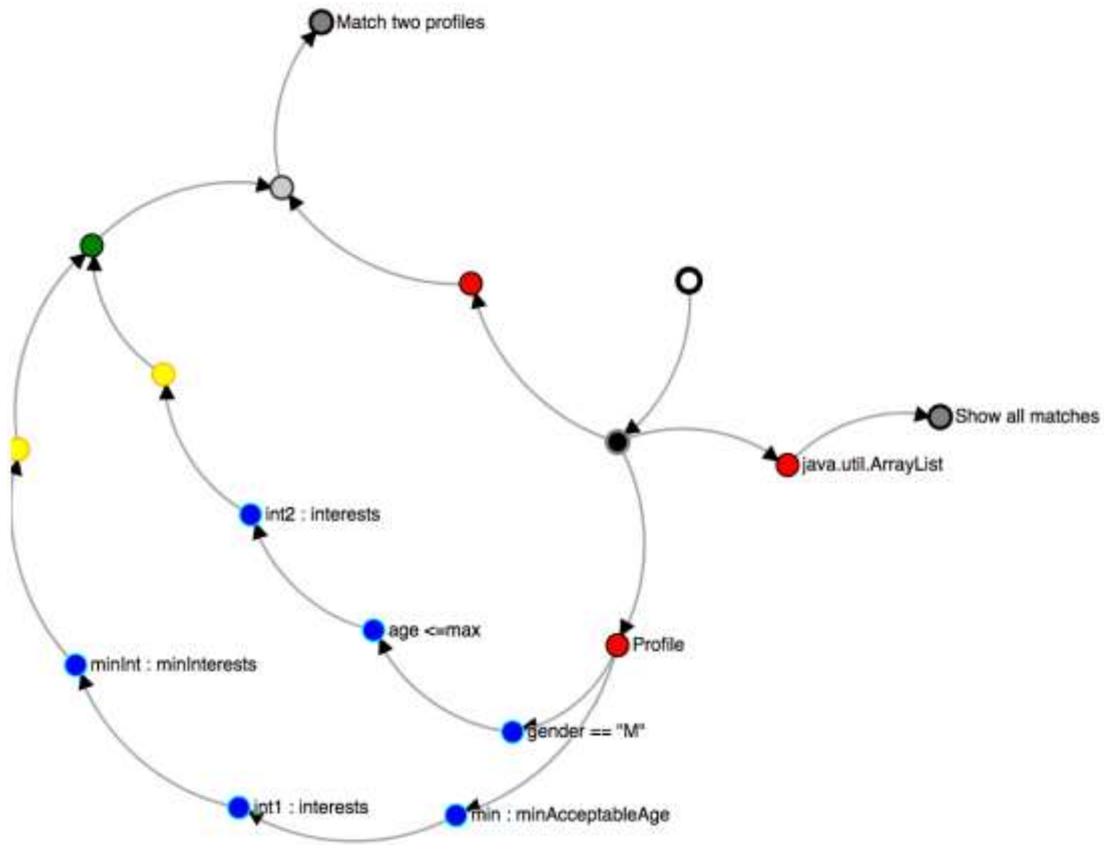

- Data driven productions - Rete
  Underutilized.
- Implementations
  Performance varies.
- Architectures
  Complex and Inefficient.
- Environments
  Cumbersome tools.

Architectures: Stateless v/s State-full

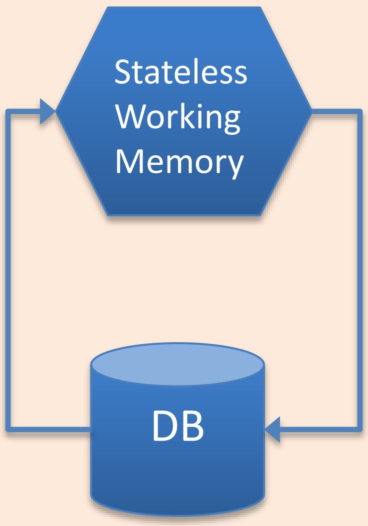

1. Create knowledge session – Rete
2. Re-create state
3. Insert new facts
4. Call ‘fireall’
5. Persist state to DB
6. Dispose knowledge session

Response time =
time to create knowledge session +
time to re-create state +
time to insert all facts +
time to fire all rules +
time to persist state to the DB +
time to dispose session

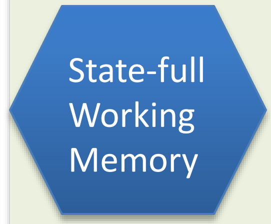

1. Create knowledge session – Rete
2. Insert new facts
3. Call ‘fireall’
4. Go to 2

Response time =
time to create knowledge session +
time to insert all facts +
time to fire all rules

Working Memory: Agenda Control & Transparency

Rete | Agenda | Tensor Flow | User Guide

Fire N | Fire Next | Fire All | Reset WM

MAIN - Create Statistical Variables (salience: 1000)
MAIN - Create Domains (salience: 1000)
MAIN - Create Categories (salience: 900)
MAIN - Compute Click through rate (salience: 1)
MAIN - Compute Click through rate (salience: 1)
MAIN - Compute Click through rate (salience: 1)

WM Statistics:

Collect Mean Square from source and Statistical Variable: 15092
Collect data from source and Statistical Variable: 15092
Populate Domains: 3773
Compute conversion rate: 2778
Compute cost to order ratio: 2450
Compute cost to revenue ratio: 2449

Console | WM Contents | Search Facts

Fact Count(before): 0, Fact Count(after): 14463
ObjectInserted: 14463, ObjectUpdated: 0, ObjectRetracted: 0

Console | WM Contents | Search Facts

Declare Type: ReportEntry   Total Facts: 1428   Download

CHANNEL	BRAND	CATEGORY	CLASSIFICATION	WEIGHT
Retargeting	Milwaukee	costOfSale	LOW	0.75
Retargeting	Milwaukee	conversionRate	LOW	0
Retargeting	Milwaukee	clickThruRate	MEDIUM	0.25
Retargeting	Milwaukee	costOfOrder	MEDIUM	0.25
Retargeting	Dremel	costOfSale	LOW	0.75
Retargeting	Dremel	conversionRate	LOW	0
Retargeting	Dremel	clickThruRate	MEDIUM	0.25
Retargeting	Dremel	costOfOrder	MEDIUM	0.25
Retargeting	Rheem	costOfSale	MEDIUM	0.25
Retargeting	Rheem	conversionRate	LOW	0
Retargeting	Rheem	clickThruRate	MEDIUM	0.25
Retargeting	Rheem	costOfOrder	MEDIUM	0.25

Learning models:

Machine Learning and Cognition

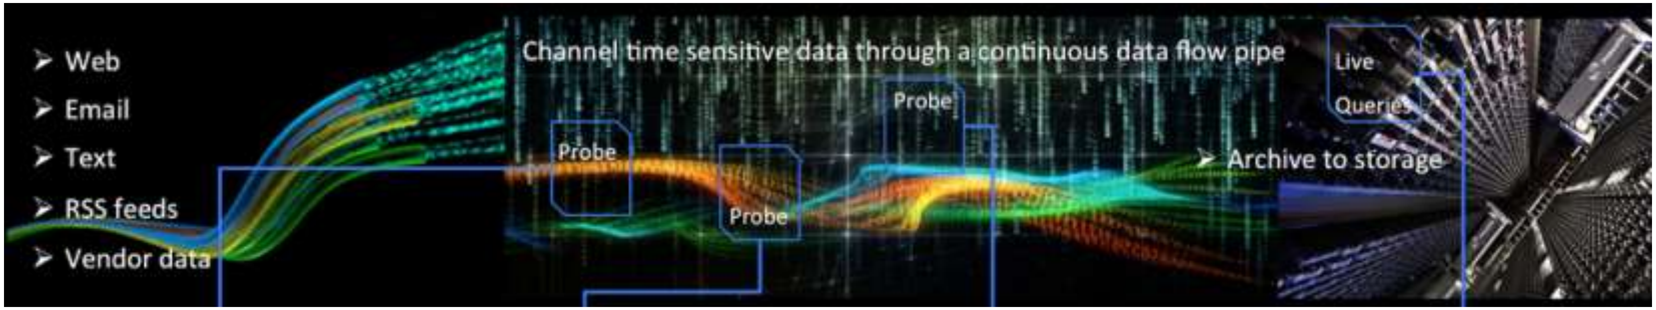

Supervised and Unsupervised Learning

- "Store-first-process-later" does not work.
- The traditional big-data analytics is too slow, and requires custom integration with decision engines.
- Deep-learning is not the answer to all ML problems.

Solution:
Adaptive Decisions Framework

Reasoning and Inference ⟷ **Continuous Learning**

Data driven productions in a state-full knowledge session

- Rule engines
- Intelligent agents
- Deductions and inferences
- Data patterns and events

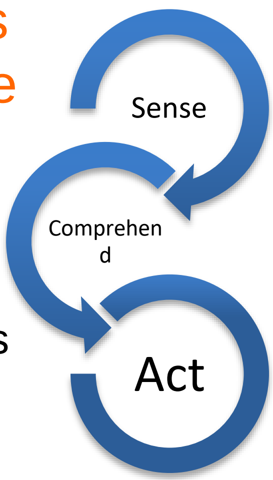

Incremental model building

- Statistical classification
- Decision tree
- Clusters
- Regression
- Neural networks
  - Resilient back propagation
  - Recurrent
  - Convolutional

Incremental Bayesian Learner

- Rules manage learning categories
- Incrementally add sample records
- Real-time update of probabilities
- Learner operates in dual mode
  - Continuous learning
  - Inference/Prediction

Learned decision tree transformed into Rete

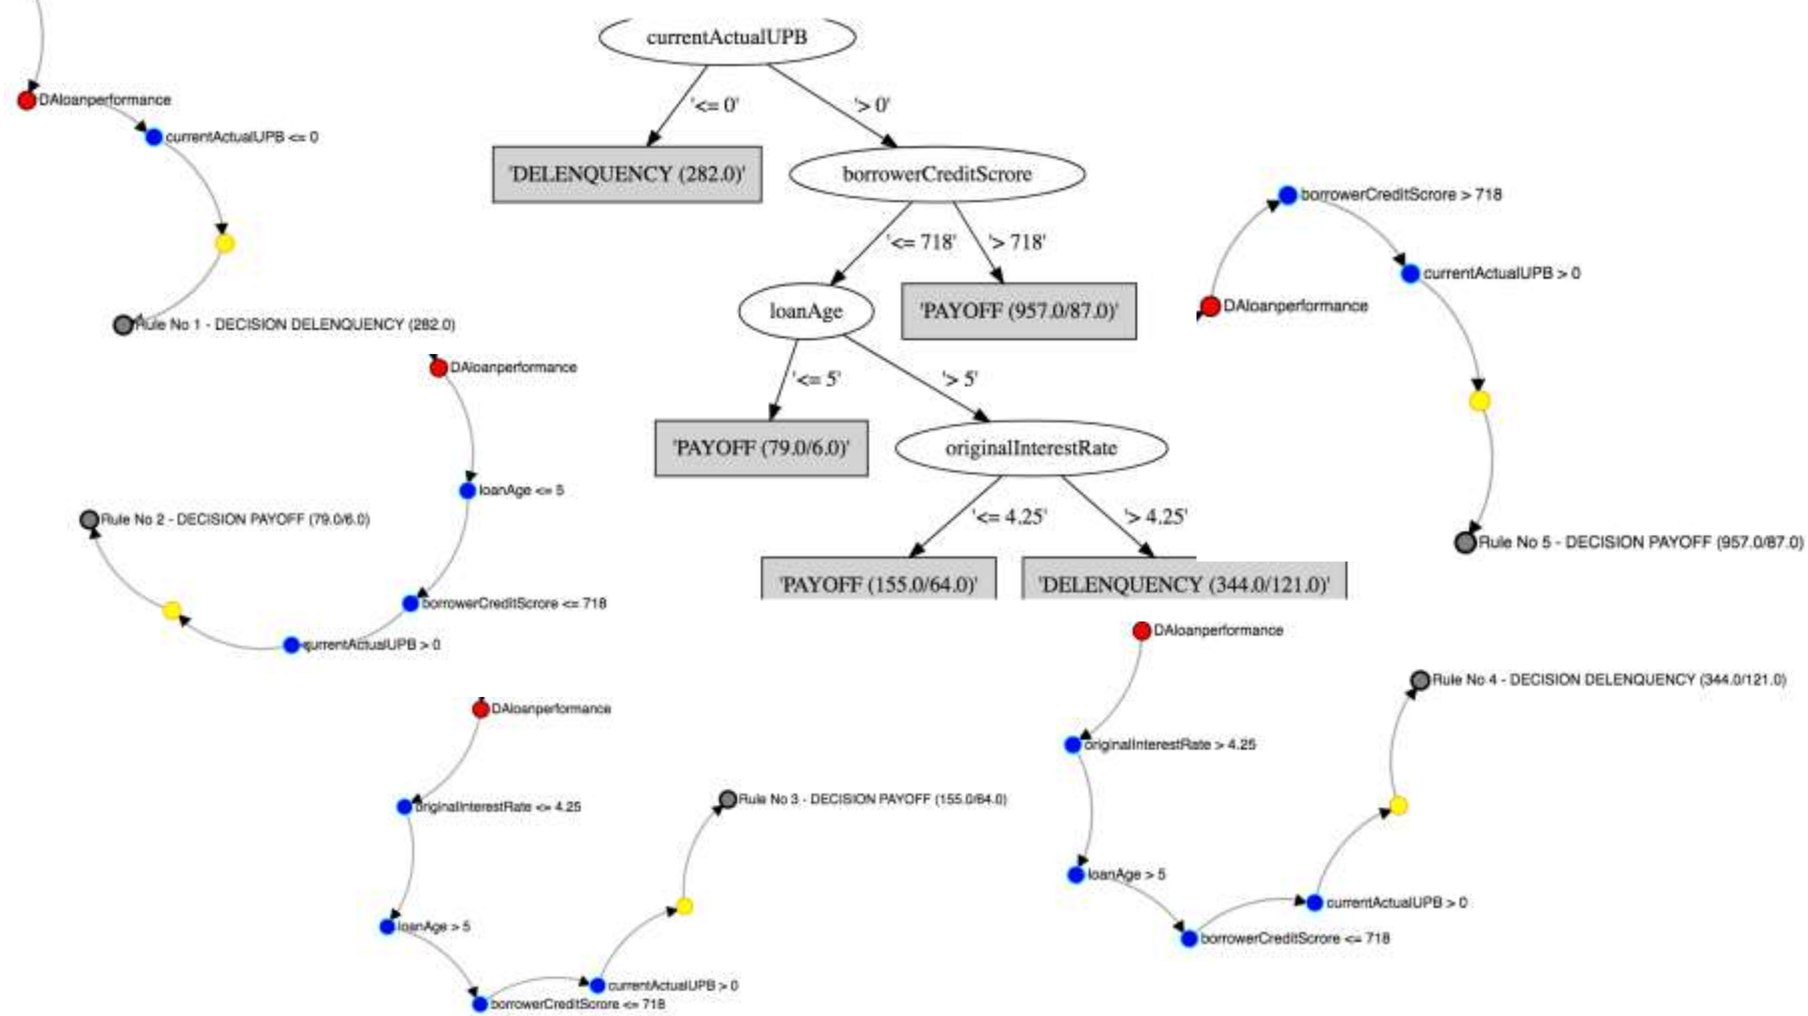

Models injected into Rete

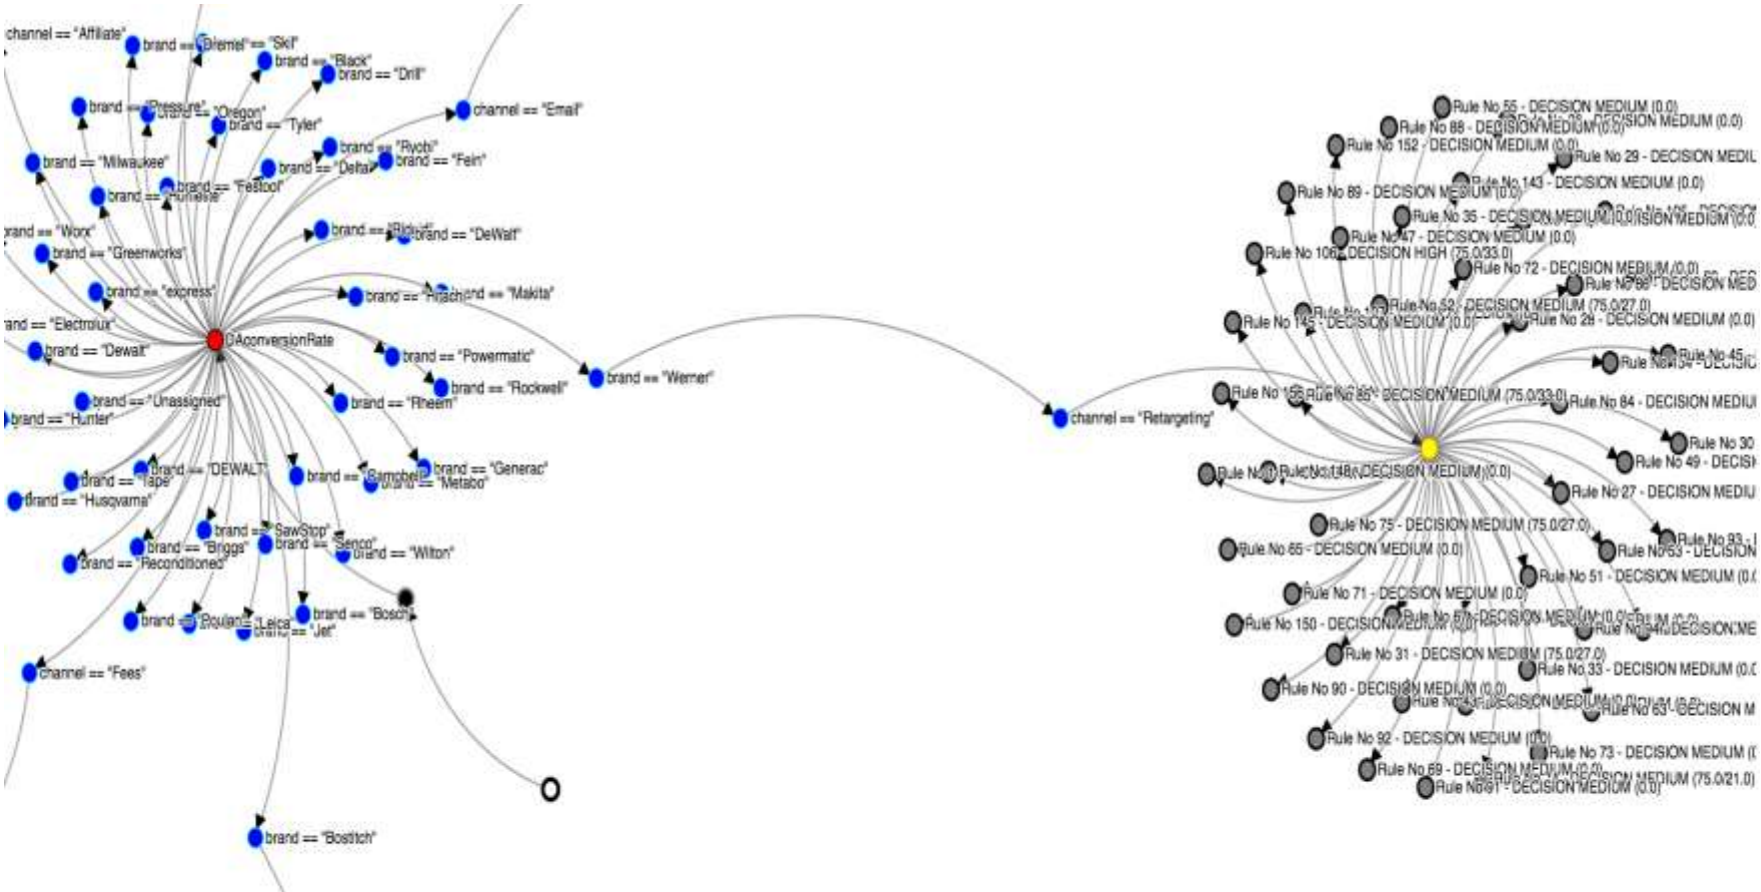

Case Study 1:
Media Purchase Decisions

CHANNEL PERFORMANCE					
|< << 2016 Week 15 >> >|					
Top 5			Bottom 5		
Rank	Channel	Aggregate Weight	Rank	Channel	Aggregate Weight
1	EMAIL	0.40	25	DISPLAY	0.010
2	INTERNAL	0.33	24	SOCIAL NETWORKS	0.010
3	NATURAL SEARCH	0.29	23	AFFILIATE	0.015
4	PAID SEARCH	0.21	22	DIRECT	0.017
5	REFERRING DOMAINS	0.14	21	PLA	0.019

1. Channel attribution (A)
2. Channel combination probabilities (B)
3. Channel-Brand combination (C)

A. probability of a channel producing revenue = revenue attributed / total revenue for the period

C.

Channel	Brand	Category	Calss
RETARGETING	Milwaukee	Cost of Sale	LOW
AFFILIATE	Dremmel	Click thru Rate	HIGH
AFFILIATE	Rockwell	Cost of Order	HIGH
EMAIL	Rheem	Cost of Order	HIGH
FEES	Electrolux	Cost of Sale	LOW
FRUGAL	Wilton	Cost of Order	LOW
FRUGAL	Delta Truck Boxes	Cost of Sale	MEDIUM

B.

	Probabilities (%)			
Channel Combination	Unique Visitors	Customers	Orders	Revenue
INTERNAL	4.05	4.66	5.16	5.87
EMAIL	4.50	5.27	5.63	4.67
NATURAL SEARCH|PLA	2.90	4.54	3.50	3.60
PAID SEARCH	4.00	2.83	2.67	2.95
DIRECT|EMAIL	0.24	2.78	2.91	2.48

Statistics on KPIs and categories

Frequency distributions

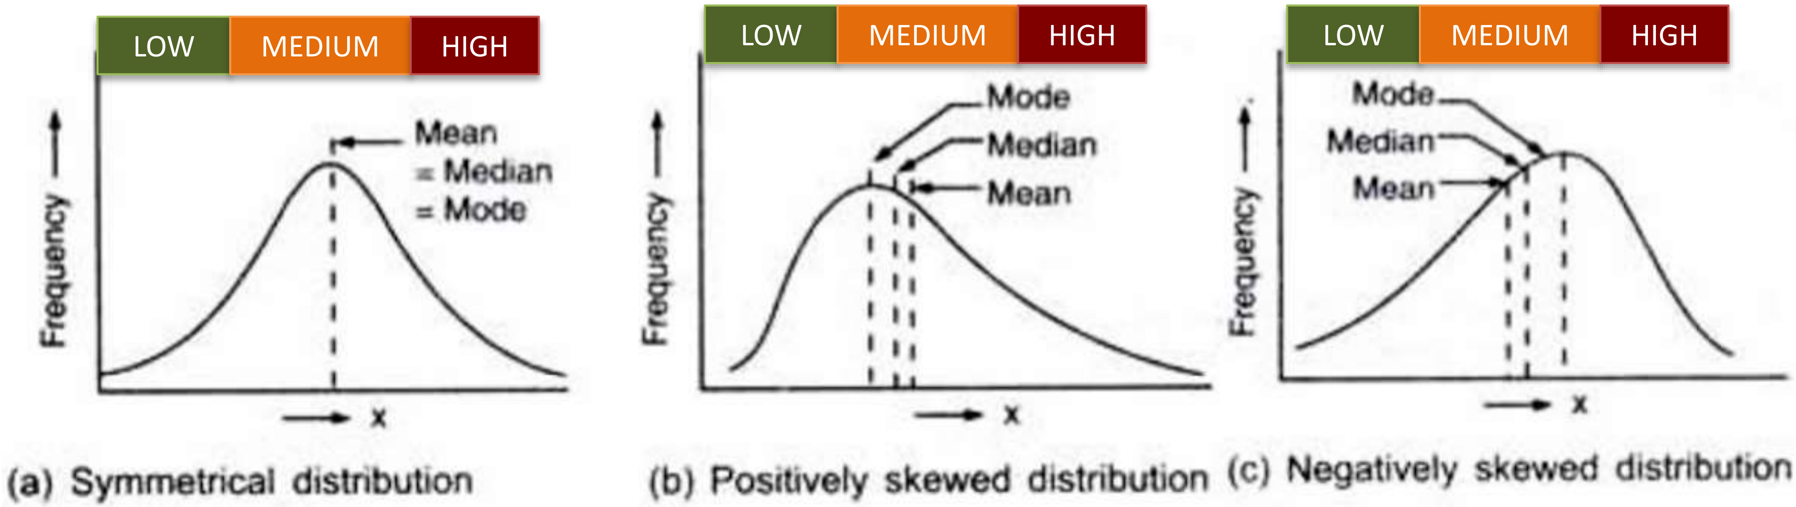

Depending on the distribution, mapping to categorical values will utilize mean and median. ***Adaptive Decision Framework*** provides Statistics template to automate this normalization process.

Results: Single retailer daily process

Daily Data Volume per Retailer		Rules count	Rules fired	Facts count		Response Time
Sales	Omniture tracking			Initial	Final	
30,184	80,000	20	373,668	115,704	127,264	793

Run	State-full	Stateless	% Improvement
1	140	360	61%
2	166	360	54%
3	66	360	82%
4	113	360	69%
5	138	360	62%
6	79	360	78%
7	68	360	81%
8	64	360	82%
9	71	360	80%
10	62	360	83%
	96.7	360	73%

Case Study 2: MBS Performance Prediction

- Predict individual loan performance
  - 25 M residential mortgages acquired over 10 year
  - 1.1 B monthly performance records over 10 year
- Cluster loans and form MBS pools
  - Interest rate distance
  - Origination distance
- Simulate MBS pool performance
  - Random loan pools based on clusters
  - Monte Carlo simulation of MBS yield

Performance

Benchmark: 2 Core 24 GB RAM VM with 100 GB hard disk.

- 45 Quarters of loan acquisition and performance data since Q1 2000:
  - Stream loan data form published web site, and
  - Insert facts into state-full working memory

  4.5 hours. (Stateless: 16.75 hours estimated)
- Add sample records to learning category: 1.35 minutes/Million records. Training time:

  6 hours. (Stateless: 22.3 hours estimated)

Key takeaway

- State-full working memory enables *real-time decisions*.
- Combining decision engine with ML facilitates *adaptive/autonomous* behavior.

Conclusion

- Architecture using state-full working memory wins
- Continuous learning within decision engine enables adaptive behavior
- Performance of adaptive decision framework meets the needs of Digital Economy.
  - Volume
  - Velocity