Sabre

# Building Personalization through Reinforcement Learning

Vidhya Balakrishnan, Sr Director, Product development

Rajeev Bellubbi, Sr. Principal Scientist, Labs

# About Me

# Vidhya Balakrishnan

## Sr. Director, Product Development, Sabre

- 20+ years of experience in managing engineering teams for Data Analytics, CX and Retailing.

## Certified Brain-based Coach*

- Certified in brain-based executive coaching
- Applying neuroscience to leadership for performance optimization
- Helping technology professionals excel in effective leadership skills

## Other passions

- Keynote Speaker on Data strategy, Customer experience in DFW data circles
- Active with Women in Technology organizations
- Flip projects at home, Reading, Hiking & Traveling

* Final certification by May 2020

# About Me

## Rajeev Bellubbi

### Sr. Principal Scientist, Sabre Labs and Product Strategy

- 20+ years of experience in software development/architecture, product development and innovation.

### Travel Industry

- Worked across Sabre's broad portfolio
- Rapid prototyping and innovation
- Data driven solutioning and algorithm development

### When I am not working

- Fitness enthusiast, Tennis, Running and Traveling

How are Artificial Intelligence/Machine Learning currently impacting your day-to-day life?

Artificial Intelligence, Virtual Reality and Machine Learning will contribute $15.7 trillion to the global economy by 2030

# AI/ML: Evolution to a point of necessity

## Artificial Intelligence

Machines carrying out tasks in
a way that we consider “smart”

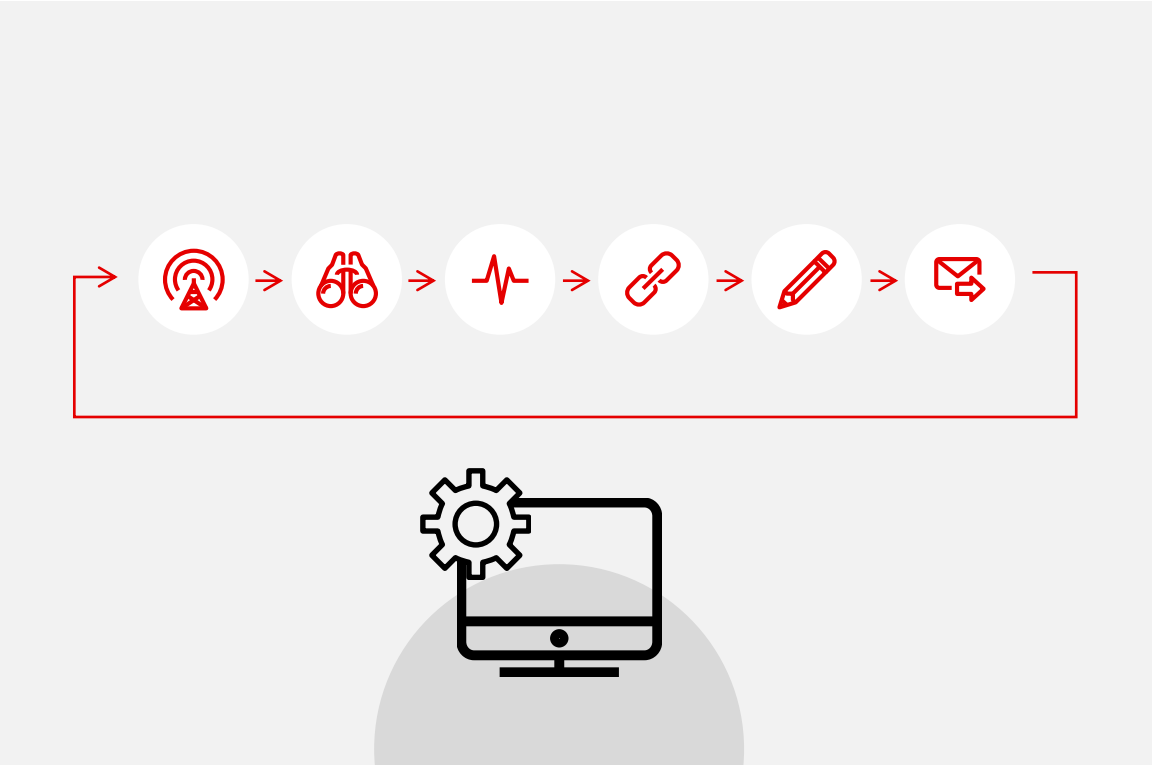

## Machine Learning

An AI application using statistical
techniques for predictive modelling

# Areas in Machine Learning

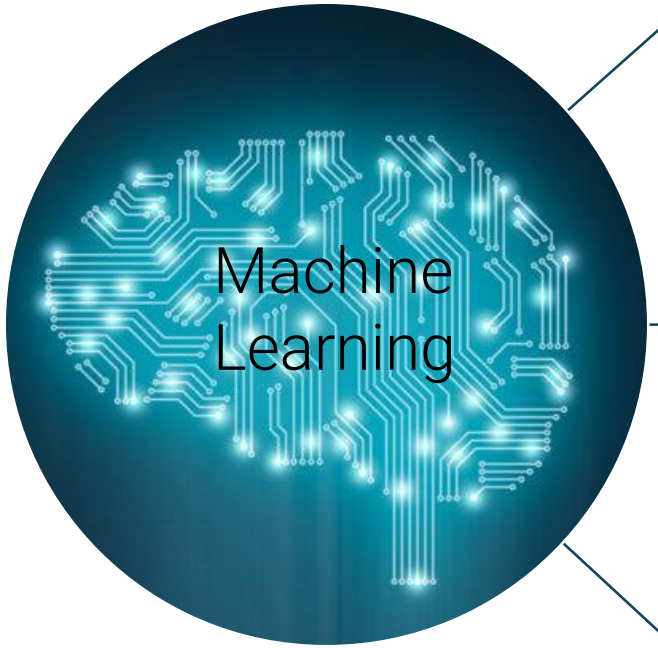


## Supervised Learning

- Classification
- Regression

- Customer Retention
- Image Classification
- Diagnostics
- Identity Fraud Detection

- Market Forecasting
- Weather Forecasting
- Estimating life expectancy
- Advertising Popularity Prediction

## Unsupervised Learning

- Clustering
- Dimensionality Reduction

- Customer Segmentation
- Recommender Systems
- Targeted Marketing
- Anomaly detection

- Big Data Visualization
- Feature Elicitation
- Meaningful compression

## Reinforcement Learning

- Realtime decisions
- Game AI
- Robot Navigation
- Learning tasks
- Skill acquisition

# What is Reinforcement Learning?

Reinforcement Learning is a ML technique that enables an agent to learn in an interactive environment by trial and error using feedback from its own actions and experiences.

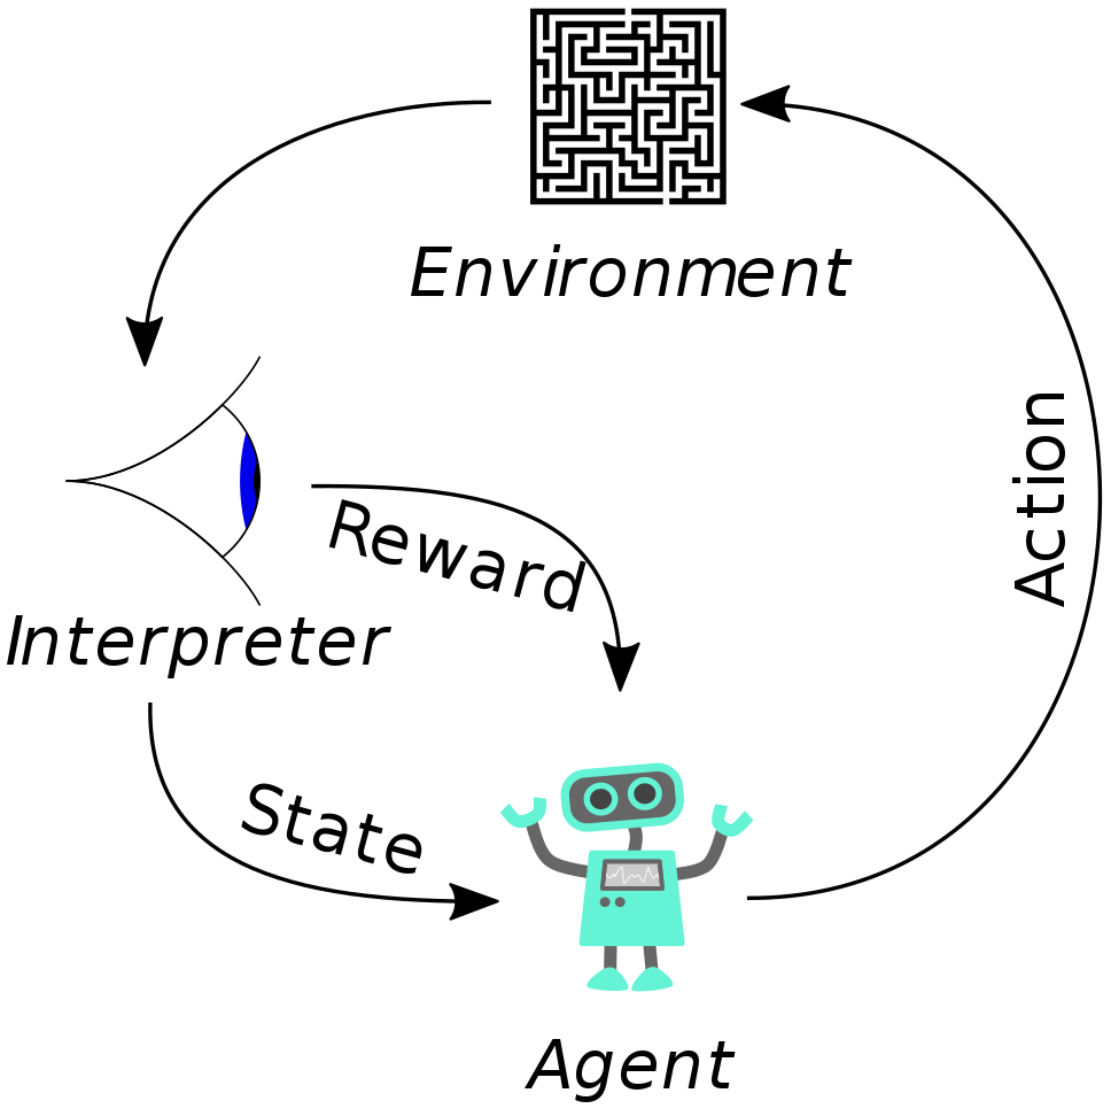


Ref: https://en.wikipedia.org/wiki/Reinforcement_learning

# Supervised, Unsupervised and Reinforcement Learning Similarities and differences

Supervised & RL

- Both use mapping between input and output
- SL - feedback provided to agent is correct set of actions for performing a task
- RL - uses rewards and punishment as signals for positive and negative behavior

Unsupervised & RL

- USL – goal is to find similarities and differences between data points
- RL – goal is to find a suitable action model that would maximize the total cumulative reward of the agent.

Reinforcement learning is based on the *Reward Hypothesis* which states that,

***"All goals can be described by the maximization of expected cumulative rewards"***

Ref: https://towardsdatascience.com/reinforcement-learning-demystified-36c39c11ec14
Ref: https://www.kdnuggets.com/2018/03/5-things-reinforcement-learning.html

Typical steps in Reinforcement Learning Algorithms

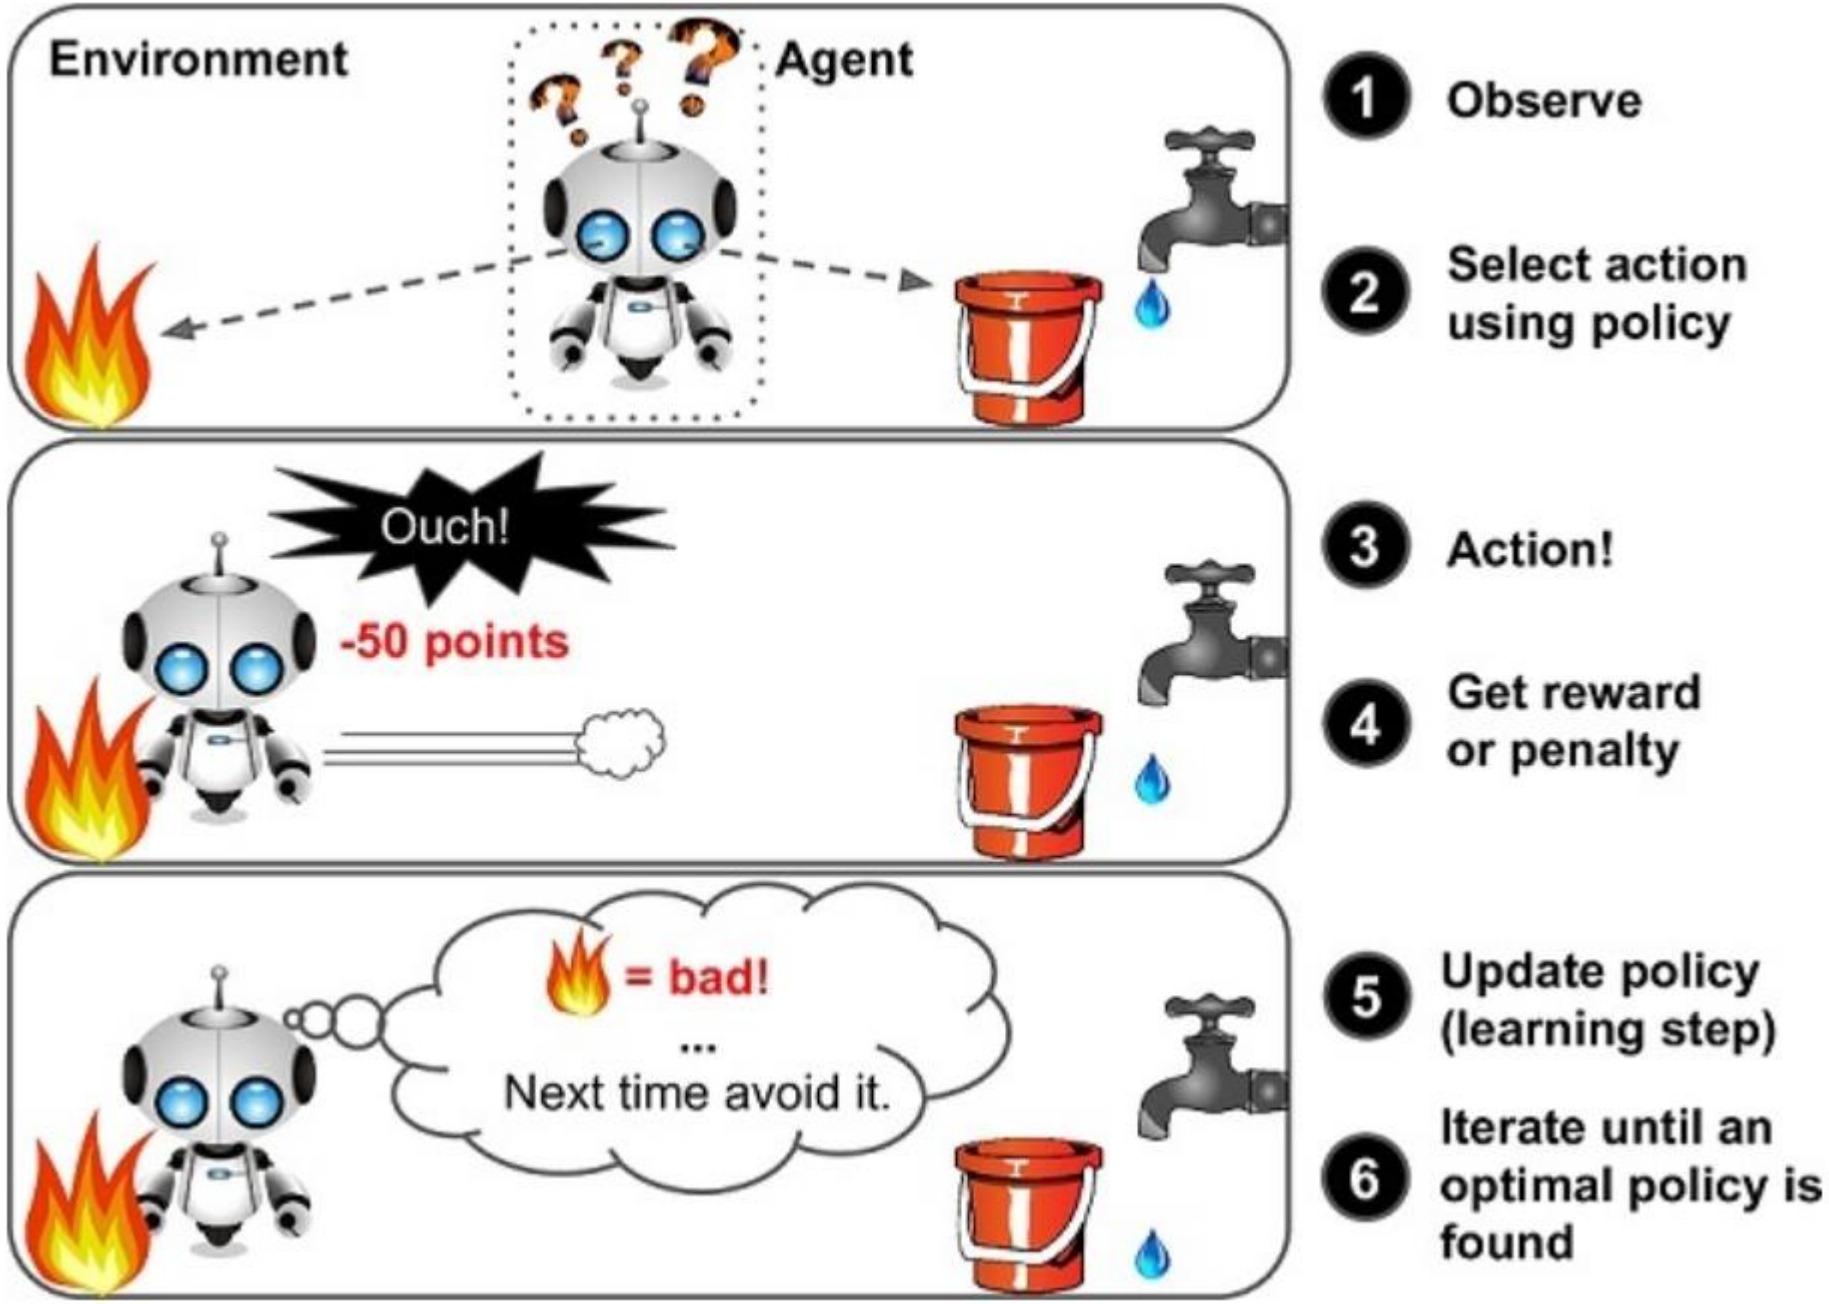


Ref: https://marutitech.com/businesses-reinforcement-learning/

# Exploration vs Exploitation trade-off

The agent must figure out the optimal way between exploring new states while maximizing its rewards. This dilemma is often referred to as the Exploration vs Exploitation trade-off

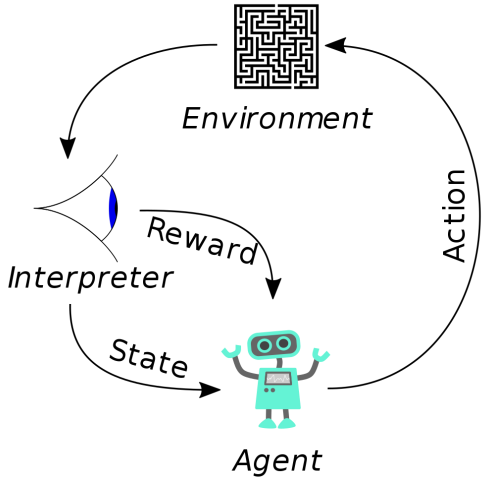


"One of the most important problems in machine learning—and life—is the exploration-exploitation dilemma. If you've found something that works, should you just keep doing it? Or is it better to try new things, knowing it could be a waste of time but also might lead to a better solution? Would you rather be a cowboy or a farmer? Start a company or run an existing one? Go steady or play the field? A midlife crisis is the yearning to explore after many years spent exploiting."
— **Pedro Domingos,** [The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World]

**Environment**: Physical world in which the agent operates

**State**: Current situation of the agent

**Reward**: Feedback from the environment

**Policy**: Method to map agent's state to actions

**Value**: Future reward that an agent would receive by taking an action in a particular state

# Common uses for Reinforcement Learning

Self driving cars

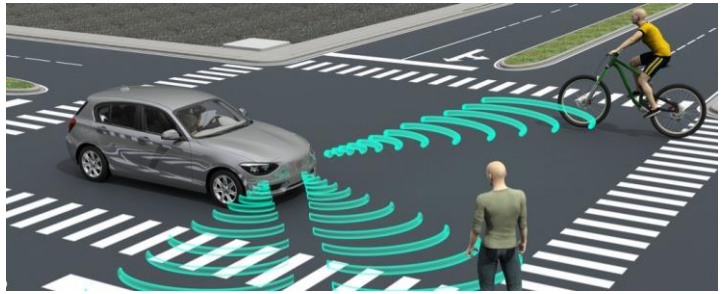

Robotic Roti Maker

$1000 for rotis

AWS Deep Racer

Google's AI AlphaZero

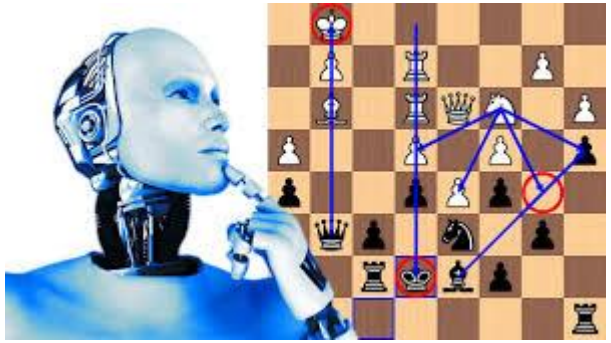

Took 4 hours to learn the moves and defeat a Grand Master

Figure: https://www.sciencealert.com/self-driving-cars-may-soon-be-making-moral-decisions-like-humans
Video: https://aws.amazon.com/deepracer/

# Business uses for Reinforcement Learning

- Optimizing space management in warehouse
- Driving financial investment decisions
- Effective treatment in medical industry
- Optimizing fleet operations for Customer Delivery
- Dynamic pricing is made possible through reinforcement learning
- eCommerce personalization

# So how is this relevant in the Travel Industry?

- Personalized Offer Management
- Dynamic Pricing

# TRAVEL INDUSTRY: RICH OPPORTUNITIES FOR AI TECHNOLOGY

A recent McKinsey study suggests the travel industry has strong future potential to apply AI technology

- Greenfield opportunities for new applications
- Next-generation versions of existing support tools
  - New and improved features
  - Higher levels of automation

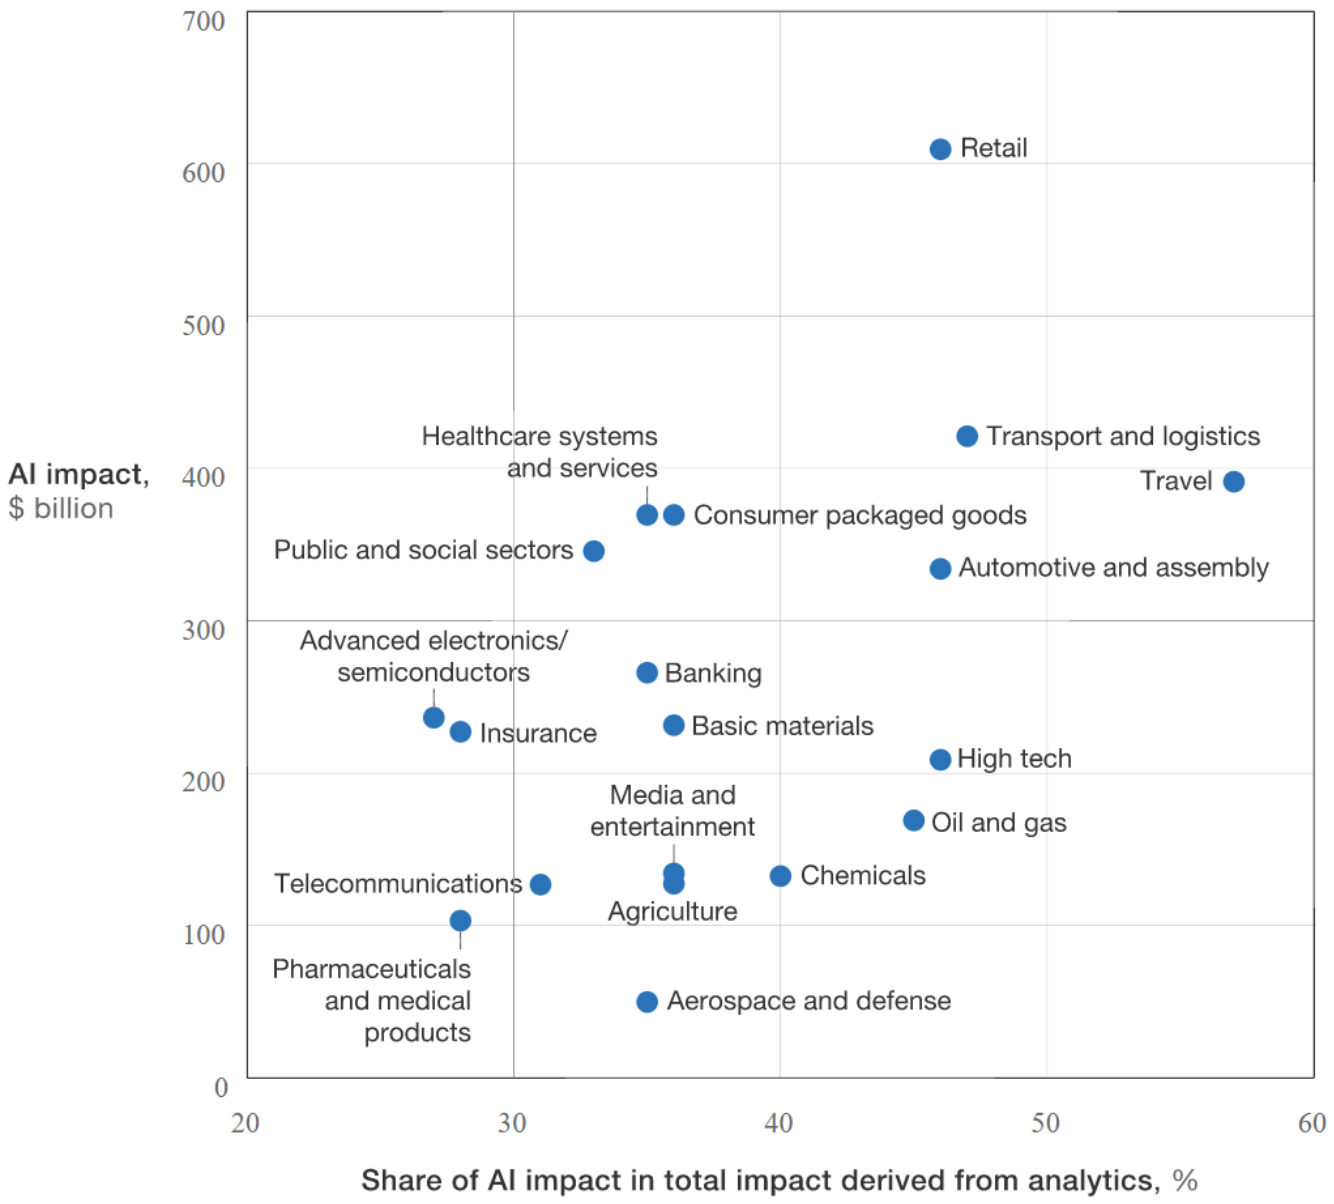

## OPERATIONS RESEARCH

Provide fast, accurate expert decision making for complex business problems that consist of a large number of variables.

## COMPUTER SCIENCE ALGORITHMS

Solve challenging business problems for complex data structures via sophisticated algorithms (e.g. delivered efficient airline shopping search results).

## STATISTICS

Enable decision making by providing exceptional insights through state of the art statistical analysis.

## MACHINE LEARNING & ARTIFICIAL INTELLIGENCE

Improve speed & quality of recommendations by enabling machines to mimic human reasoning, correction and learning.

Sabre

How is Sabre leveraging these advancements in technology?.....

Sabre

**#1 or #2**
Across all product sets

**$3.9B**
In revenue
(2018)

**SABR**
Publicly traded on NASDAQ

**~9,000**
Employees in 65 countries

**Deep**
Customer relationships

## Travel Solutions

**Travel Network**
Operating the world's largest travel marketplace, connecting travel buyers and sellers

**Airline Solutions**
Offering intelligent, connected platforms for retailing, distribution and fulfilment that help airlines realize profitable growth

## Hospitality Solutions

Providing distribution, operations and marketing solutions to the hotel industry

©2019 Sabre GLBL Inc. All rights reserved.

# Every minute of every day, Sabre technology powers travel.

**62.5M**
API Calls
via Dev Studio

**132K**
Flights searched

**3.5K**
Messages sent
to travelers

**5.2K**
Hotel rooms shopped

**$26K**
In booked
hotel revenue

**$428K**
In booked air travel

**1.1T**
SYSTEM MESSAGES PROCESSED ANNUALLY

**$120B**
GLOBAL TRAVEL SPEND PROCESSED ANNUALLY

**$1B**
INVESTED IN TECHNOLOGY ANNUALLY

# 1 BILLION+

TRIPS BOOKED ANNUALLY

Sabre

# Defining “Personalization use cases” for Travel

# ‘Experiences’ trumps ‘Products’ when it comes to competitive advantage

- Uber & Airbnb brought about a great revolution putting experiences before product
- The new trend with the GenZs is going to be on-demand, mobile first economy

https://www.briansolis.com/2016/02/26-disruptive-tech-trends-2016-2018/

# Importance of Personalization to uplift your business and revenue margins

Benefits of Personalization

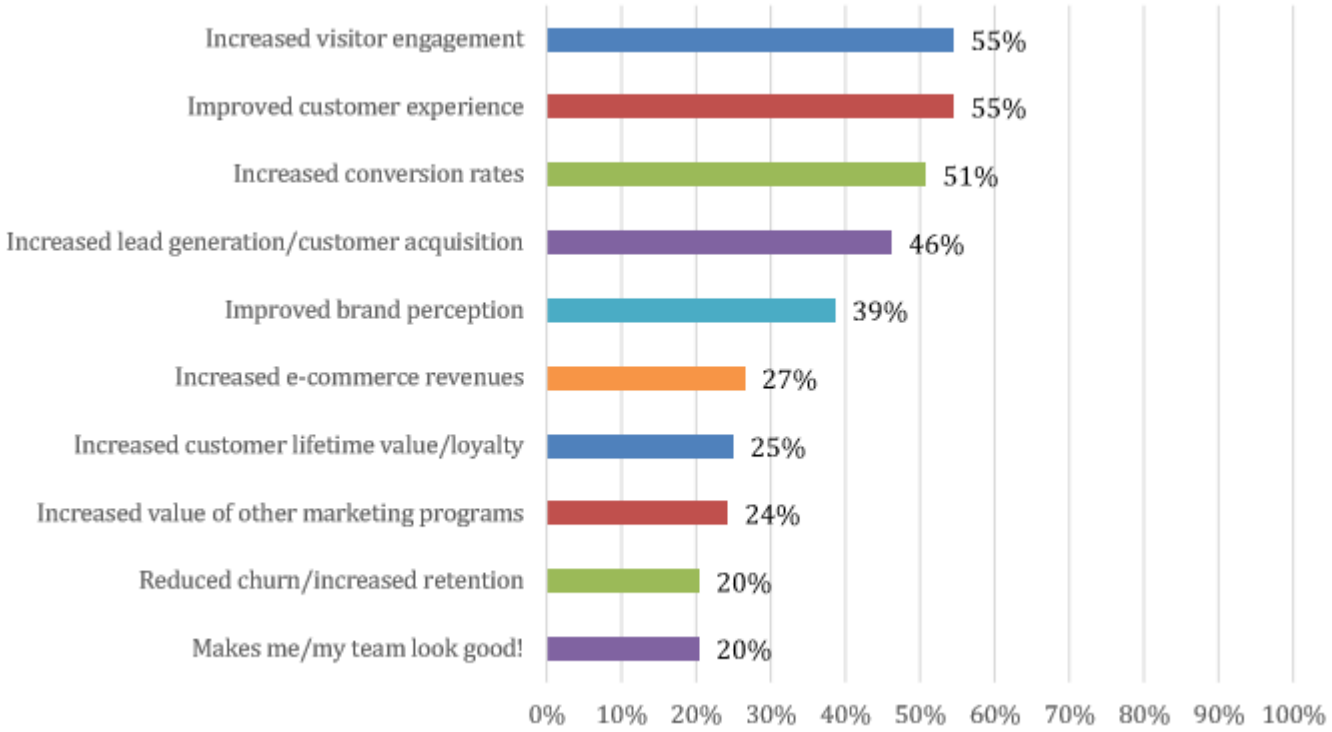


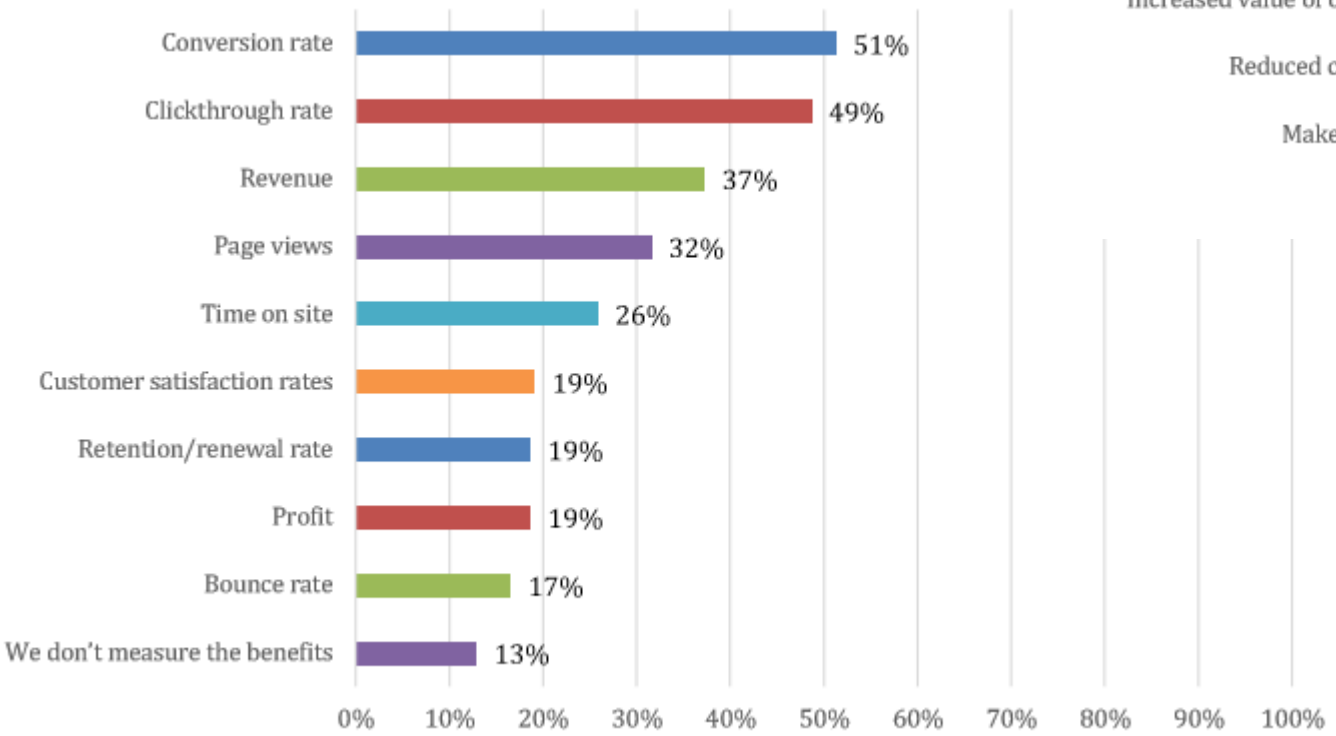


g/persona

Effectiveness of Personalization based on improved KPIs

Source: https://www.evergage.com/wp-content/uploads/2018/04/Evergage-2018-Trends-in-Personalization-Survey.pdf

# Personalized travel with Intelligent Retailing

## Travelers → Personas → Personalization

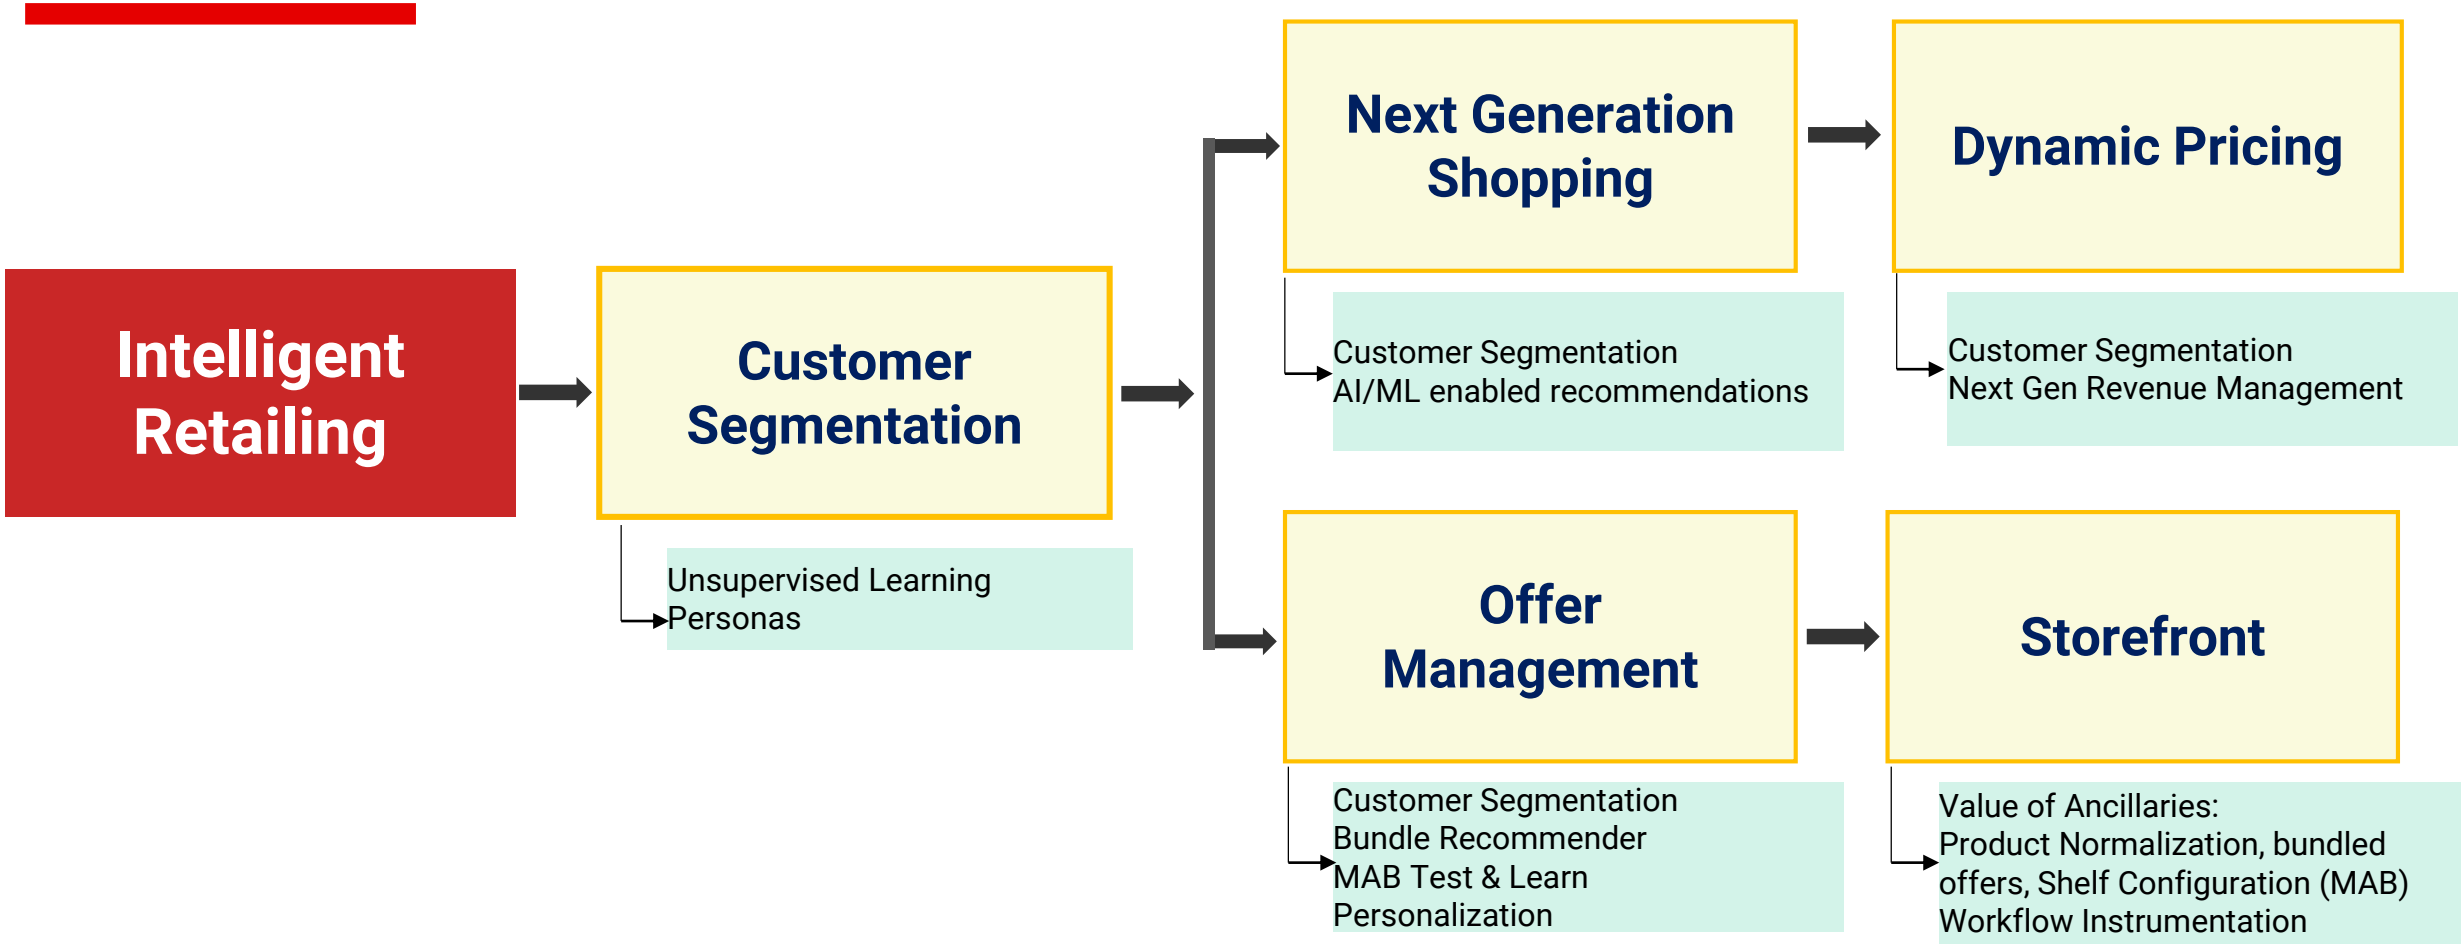

# A persona-based recommendation engine to maximize conversion rates of the air bundle – airline view

**Recommendation Engine**

Recommends bundles based on travel context

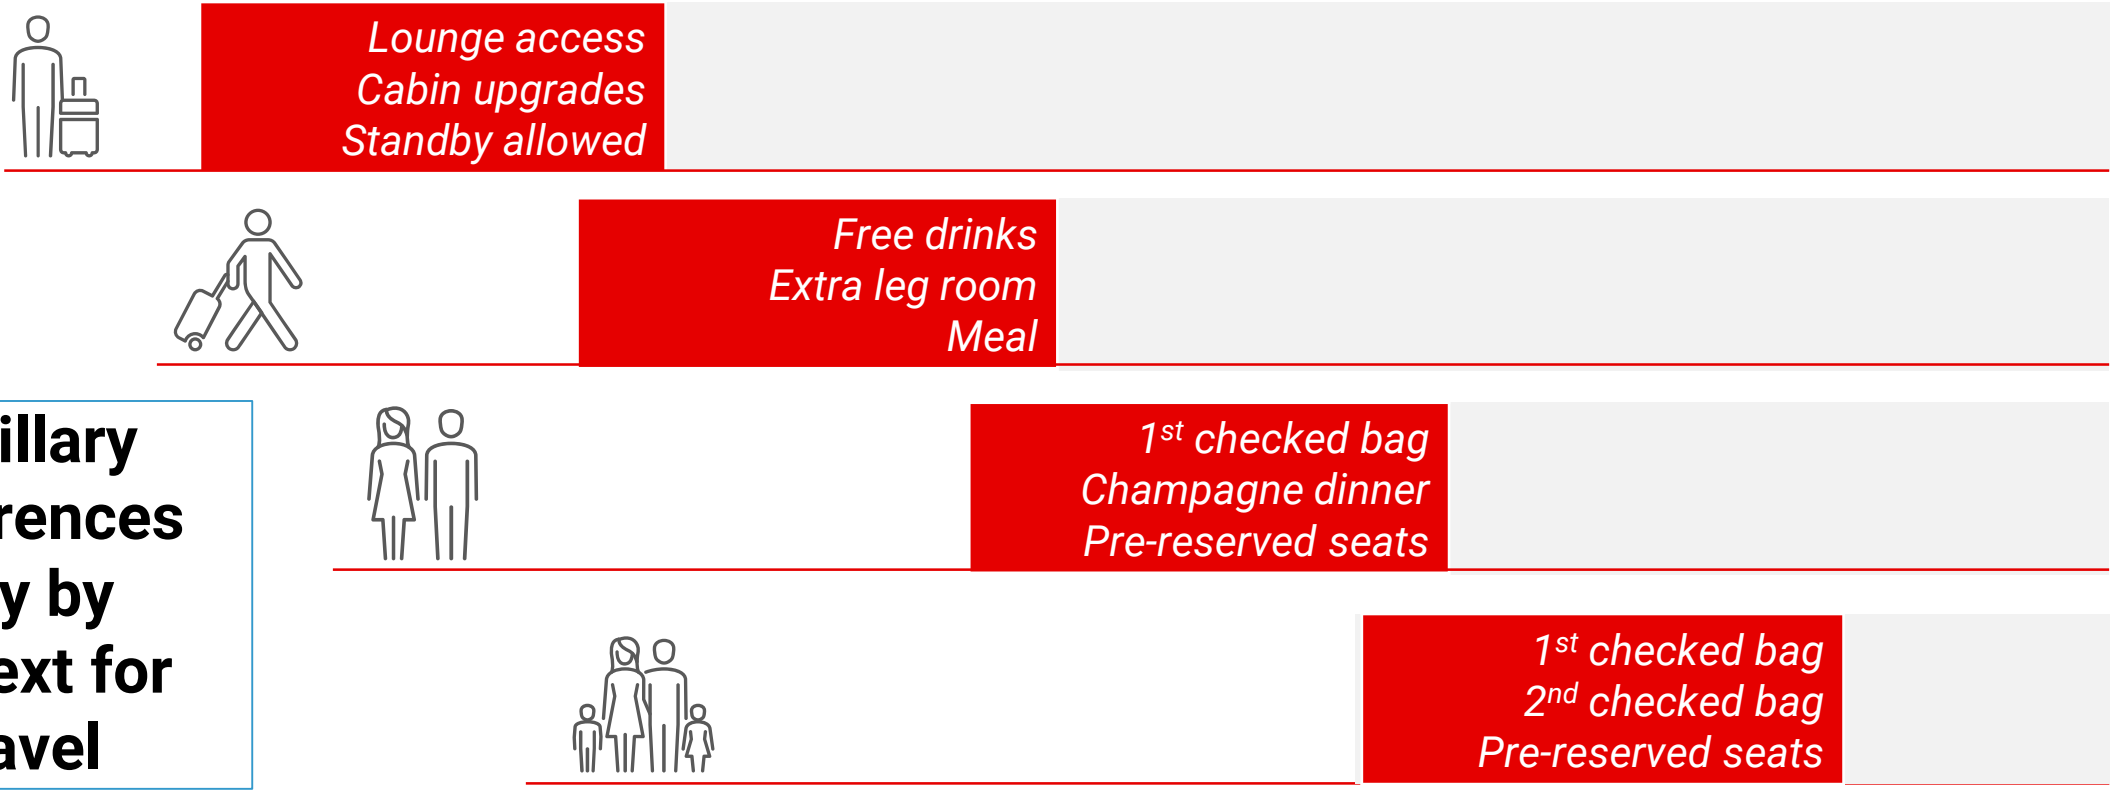


**Ancillary preferences vary by Context for Travel**

# Data & Segmentation

## Traditional segmentation

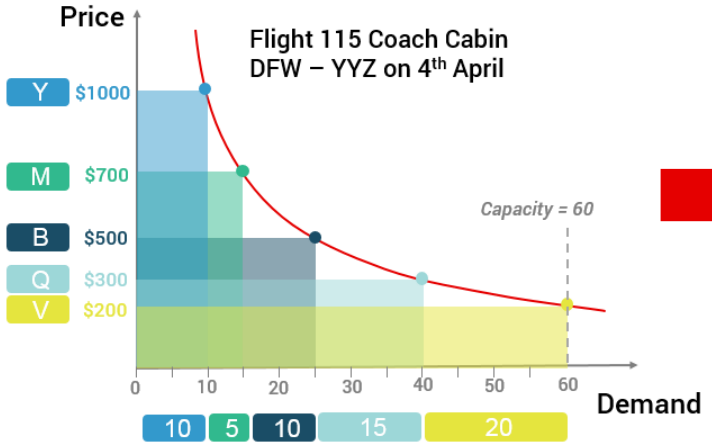


## Context-based segmentation

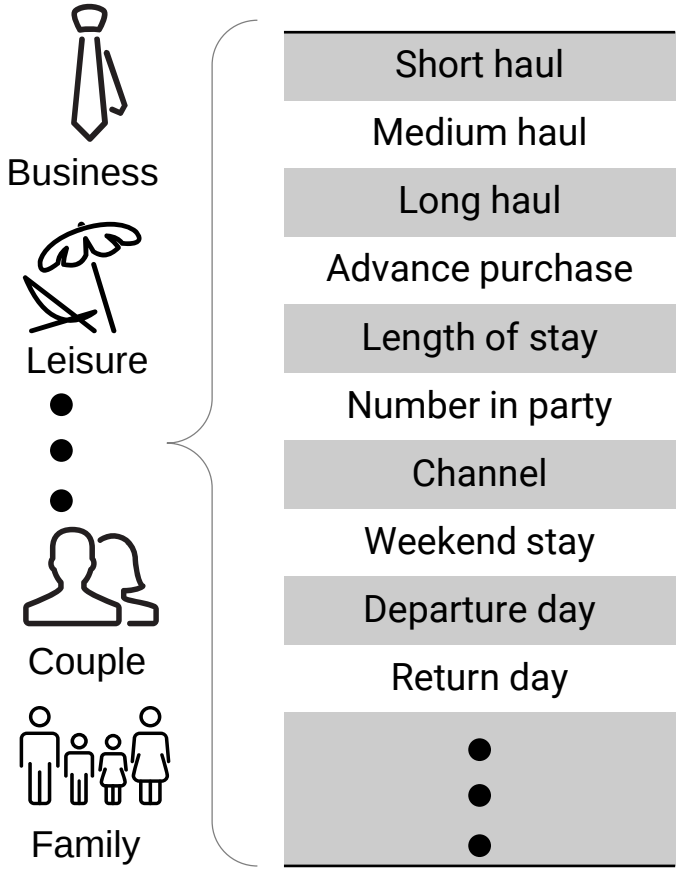


## Travel extras

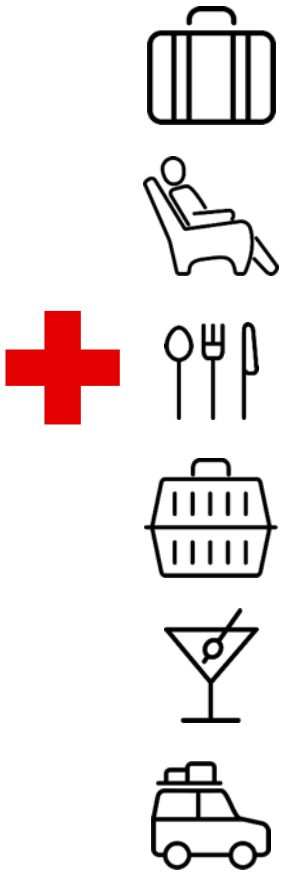

## Personalization

*Segment of ONE*

From personas to 1:1 personalization

# Advanced decision support for personalized offers

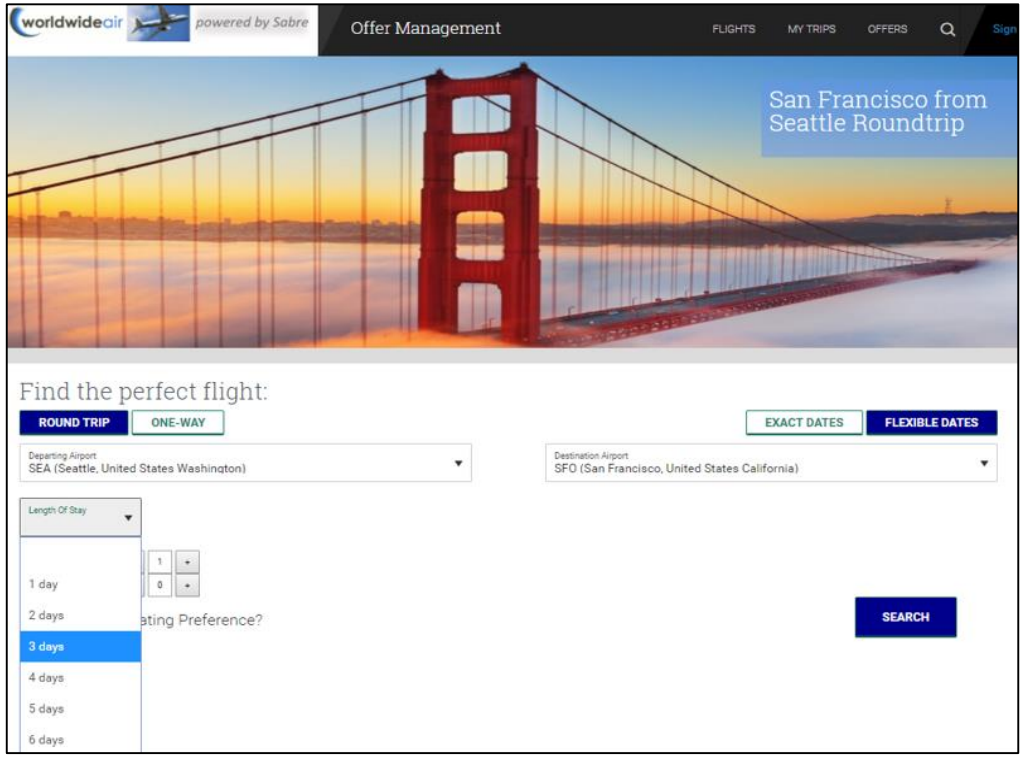


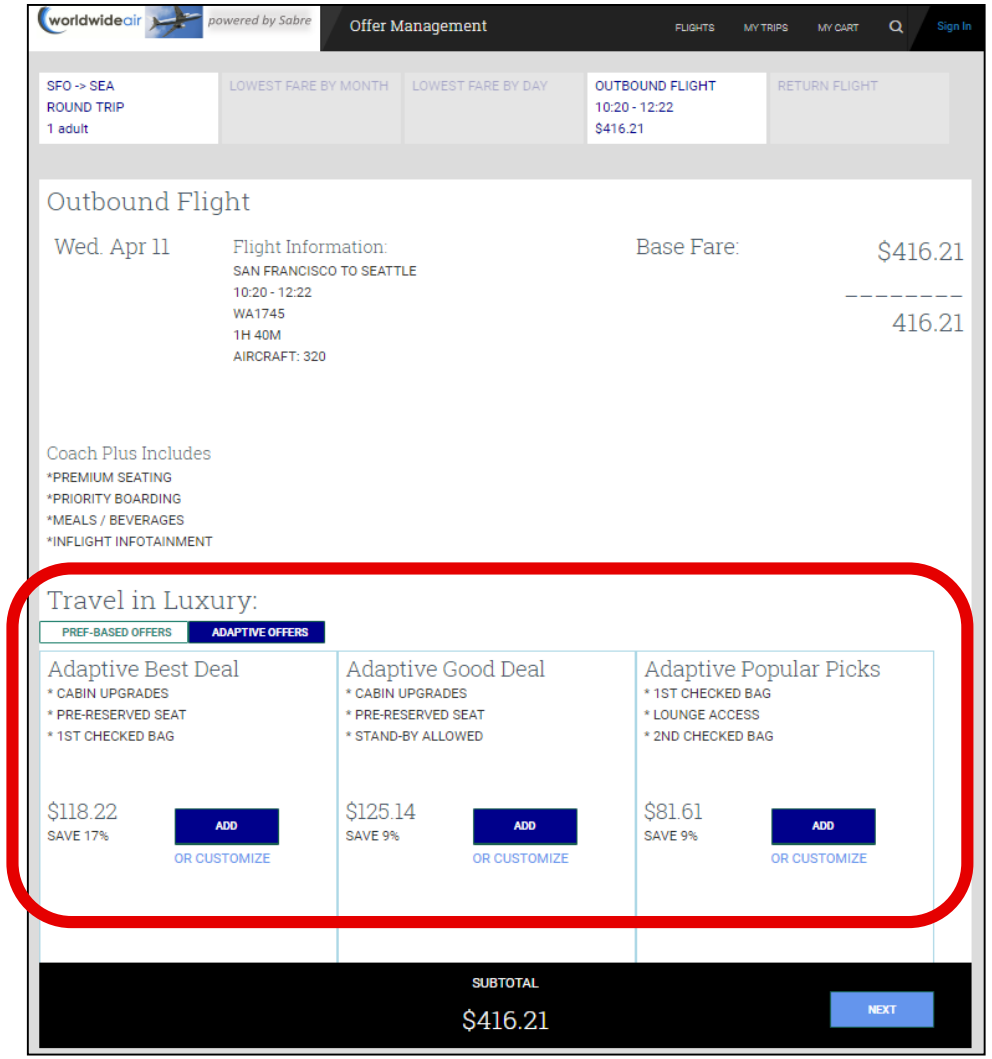

# “Smart” experiments drive intelligent recommendations for air shopping, offer management, intelligent storefront, ancillary pricing and much more…

- **Business Question:** Which option is the best alternative?
- Experiment by allocating options across requisite dimensions, (e.g., dates, day-parting, etc.)
- Use **multi-armed bandit (MAB)** technique to test the associated metrics at various alternative scenarios and continuously tune the model to focus on the best performing alternatives

*Fixed allocation*

A/B Testing

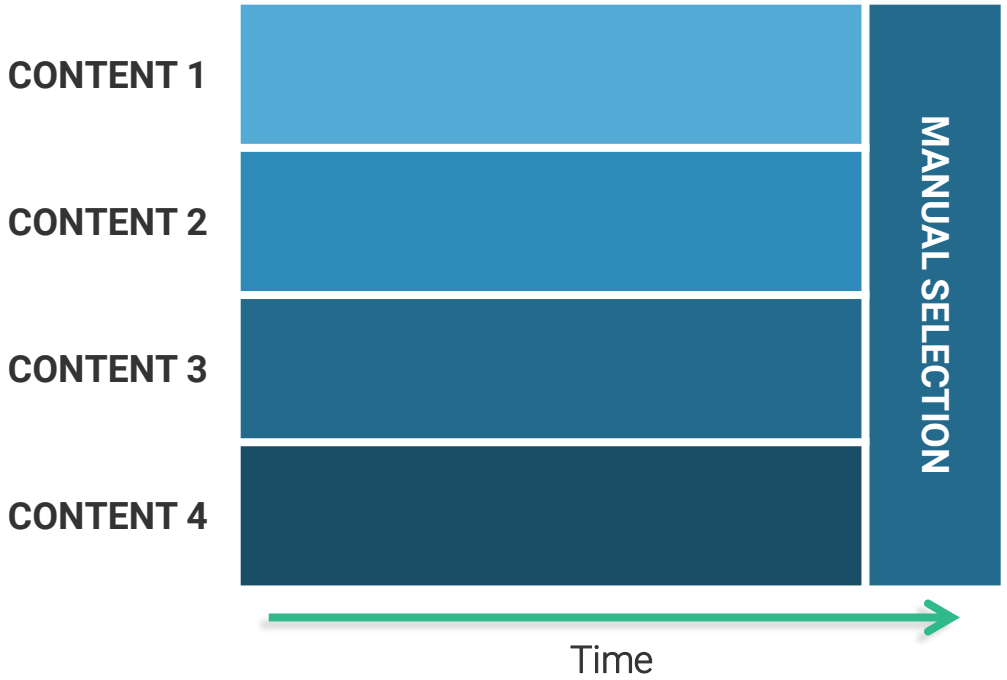

*Continuous reallocation*

MAB

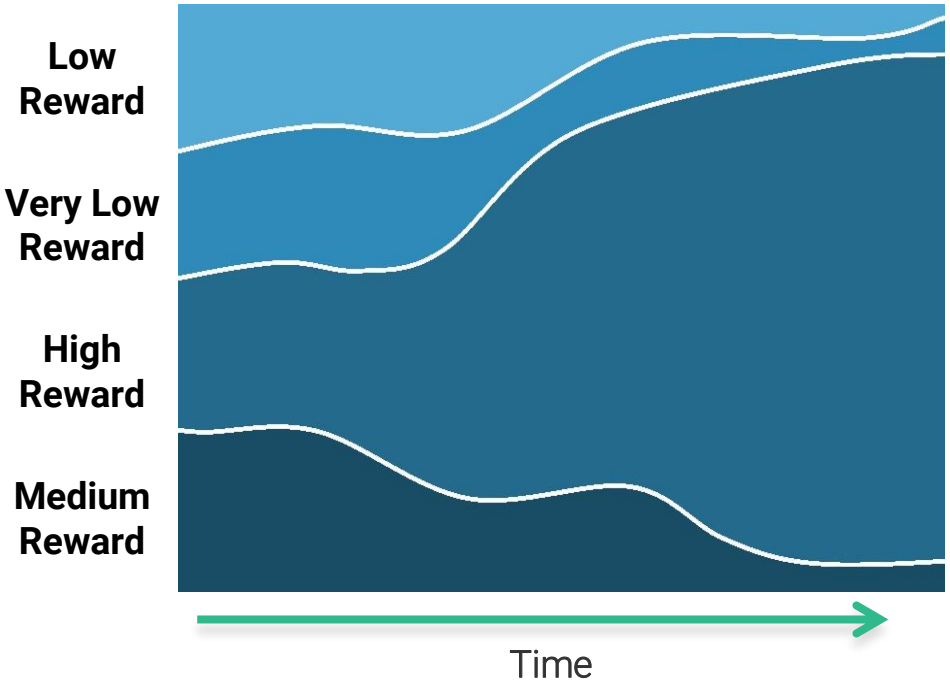


Adapted from: Frosmo, Multi-Armed Bandit, Joni Turunen, 2017

# Under the Hood of the Machine Learning Application

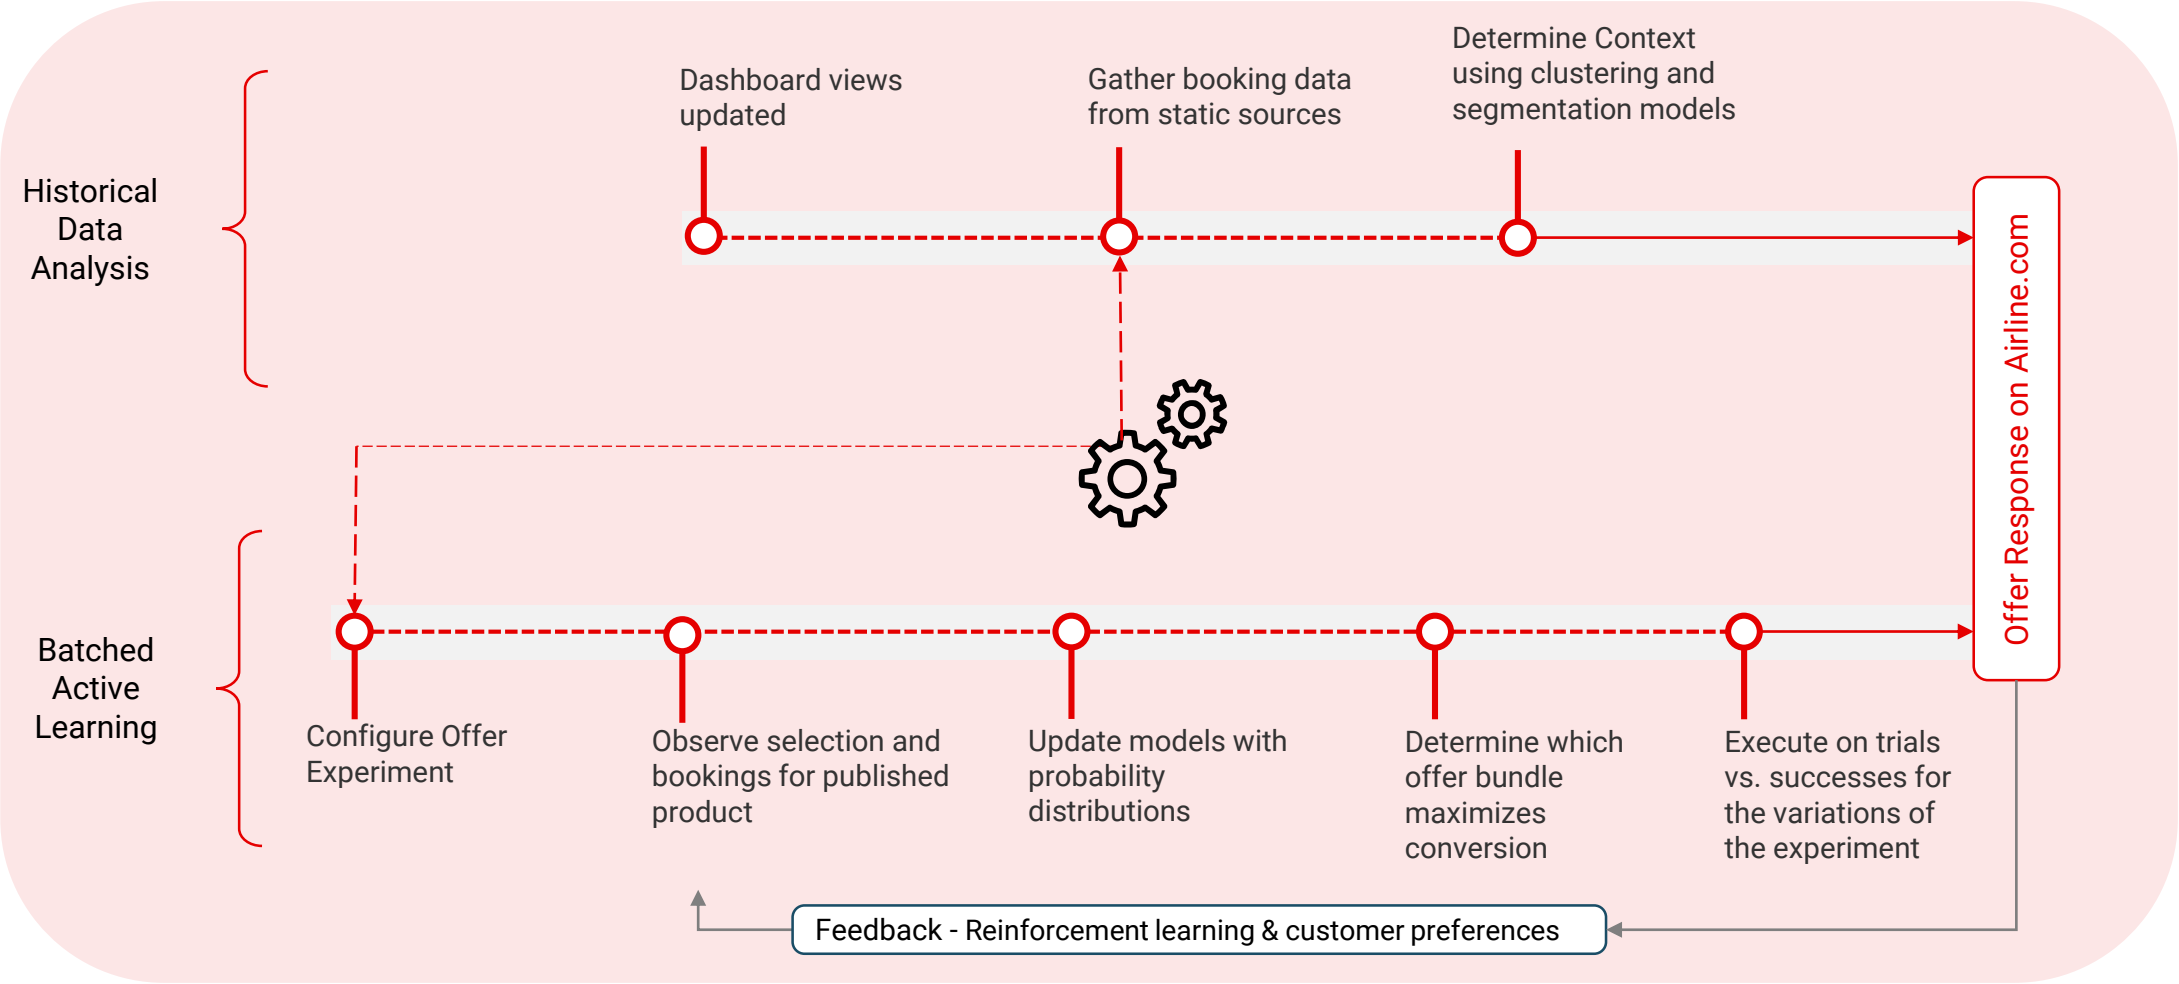

# Multi-armed bandit: best reward or more information?

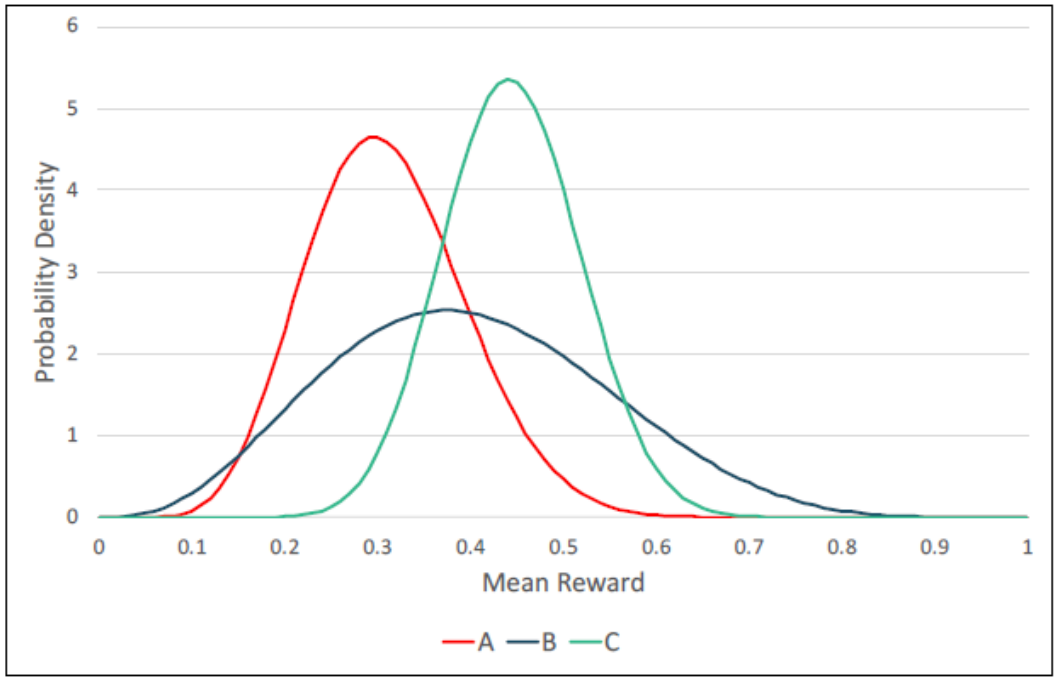


Posterior distributions of mean rewards observed in previous time periods

- Exploiting arms that yielded high rewards in the past for immediate rewards
- Exploring other arms to gather more information for long term rewards
- Develop a sequential strategy for pulling arms to balance this tradeoff and max the total expected reward or min total expected regret

# More exploration = less uncertainty

## Distribution gets more concentrated with more observations

- As we observe more activity for a given product and customer segment, we gain more information about the underlying distribution
- Conversion rates
- Variance
- Often modeled as Beta, Gamma or Lognormal

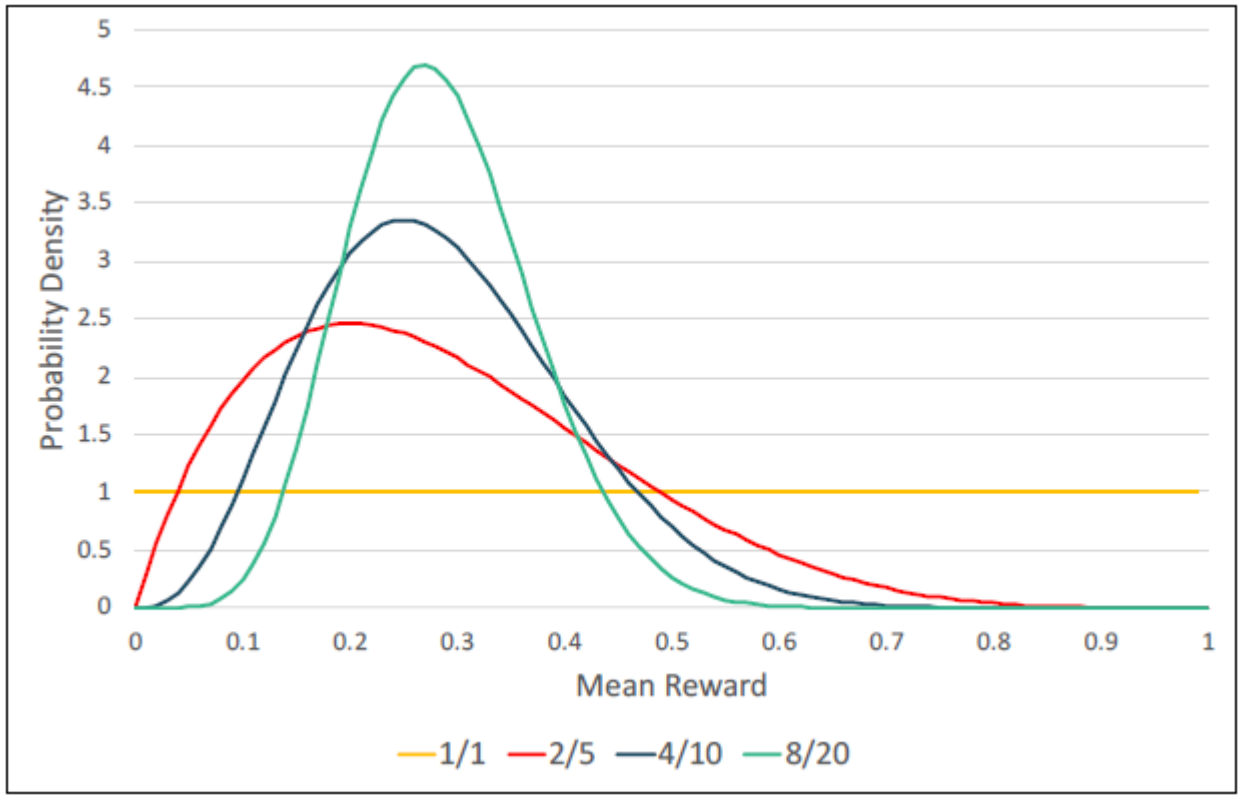

# As data are accumulated, these MAB models converge to recommended solutions by customer segment

Bundle combinations that are the most preferred for this customer segment

Note: on this item, probability dropped to zero but later increased after a few successful trials. So MAB models can automatically adapt to changing market conditions or customer preferences

WORLDSTARAIR

WorldStar Rewards

Log In

BOOK

CHECK IN

FLIGHT STATUS

MY TRIPS

ROUND TRIP

Departing from?

Where to?

Travelers
1

Cabin
Economy

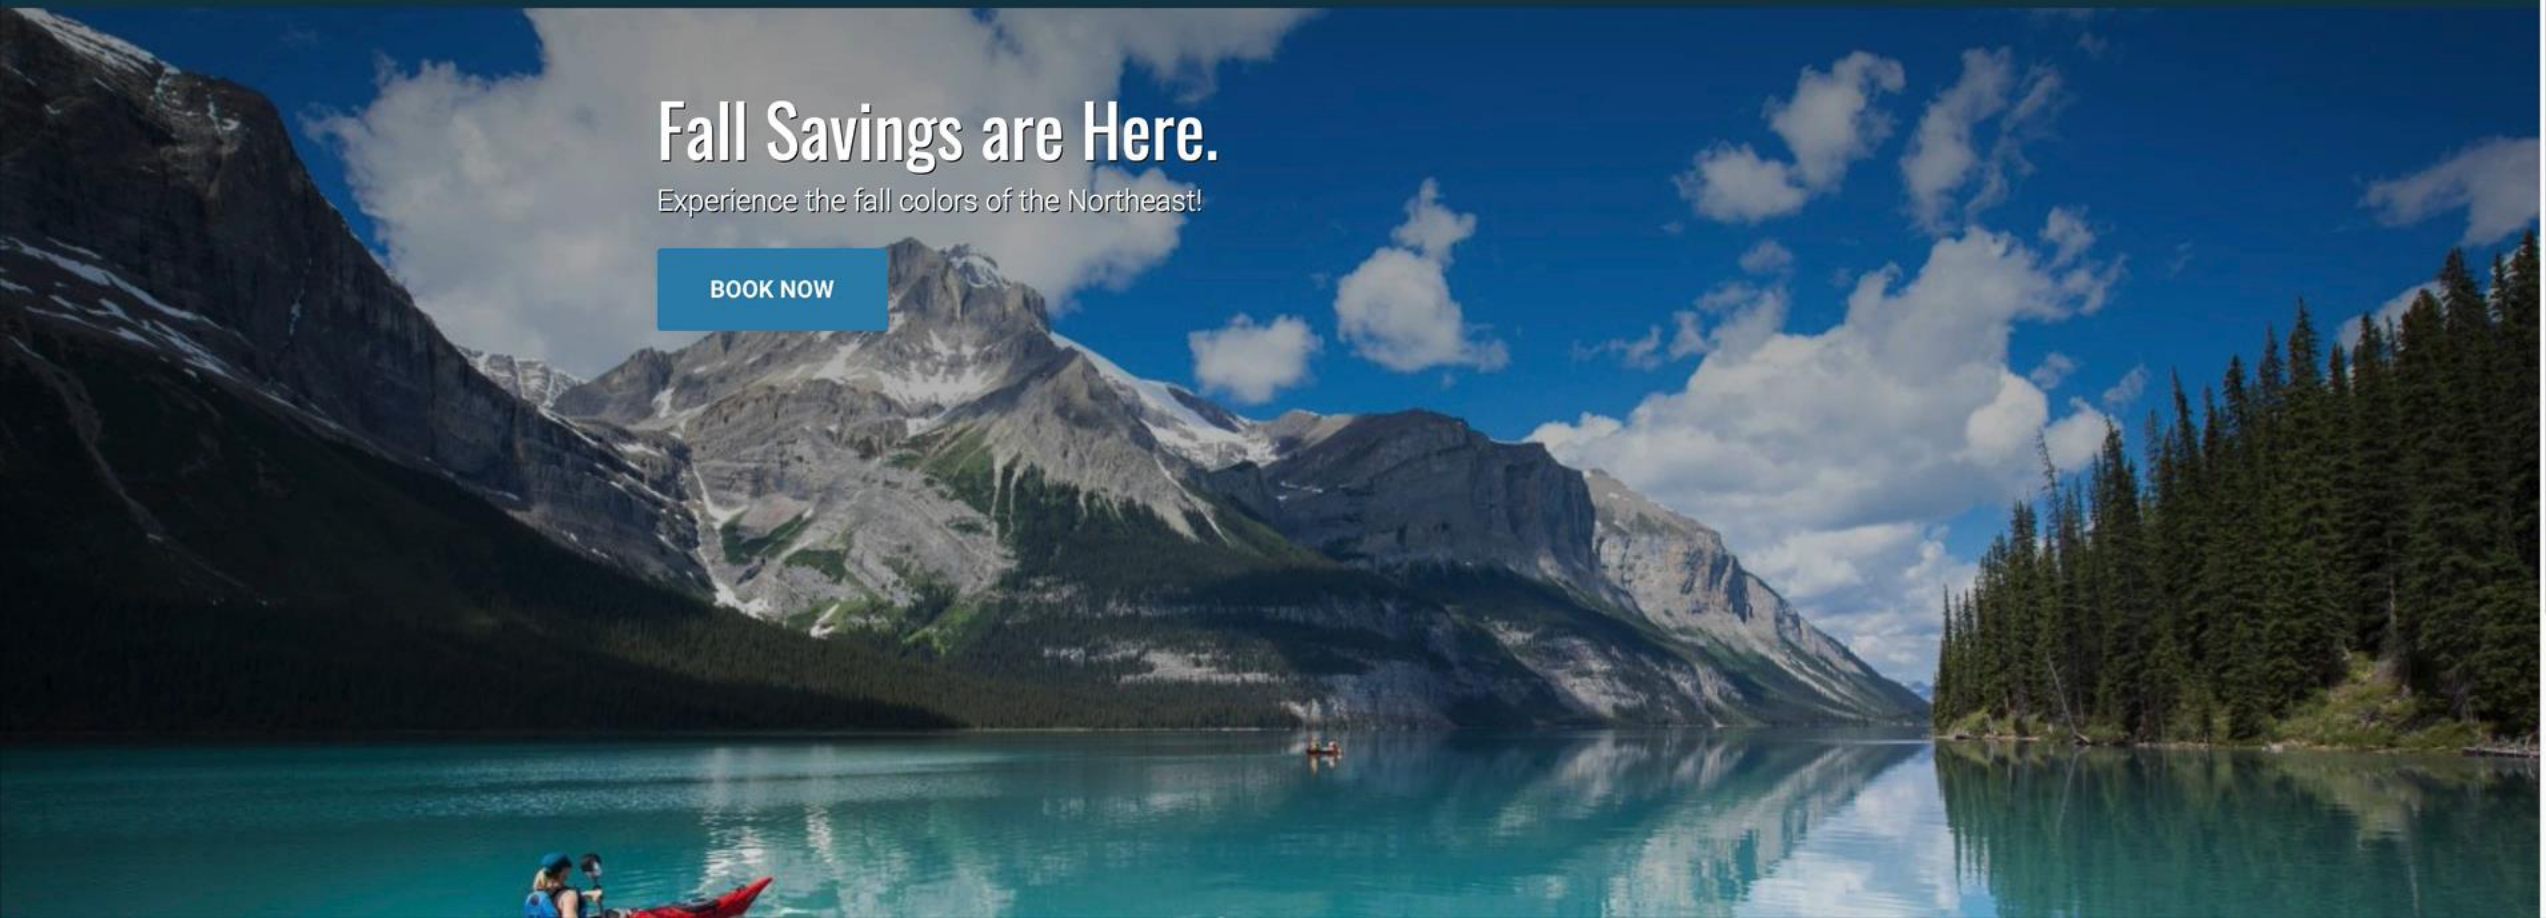

Q & A

Sabre®

# SynXis Intelligent Retailer

## Create Infinite SKU Non-Room Offer Engine that sells ancillaries/add-ons

- Sells both physical inventory (any object that a traditional shopping cart would sell such as beer in your room, slippers, robes, hotel branded coffee cups, umbrellas, etc)
- Non-physical inventory (late check out/early check in, guarantee policies, parking, internet, etc)
- Initial integration with Hospitality Solutions Services & SynXis Booking Engine, followed by Voice Agent and Property Hub

## Machine Learning Recommender Engine

- Product ordering and yielding recommendations
- Later we will add customer segmentation with DX/Google capability

# Retail recommender model scope

## Research assisted SynXis in applying AI/ML models to enable smarter retailing decision support

- Began integration testing in late 2019
- Went live in production this month

## Three main models used:

- Trip-purpose segmentation (classification rules)
- Recommender model (using multi-armed bandit reinforcement learning)
- Simulation (to validate and compare the performance of different retailing algorithms)

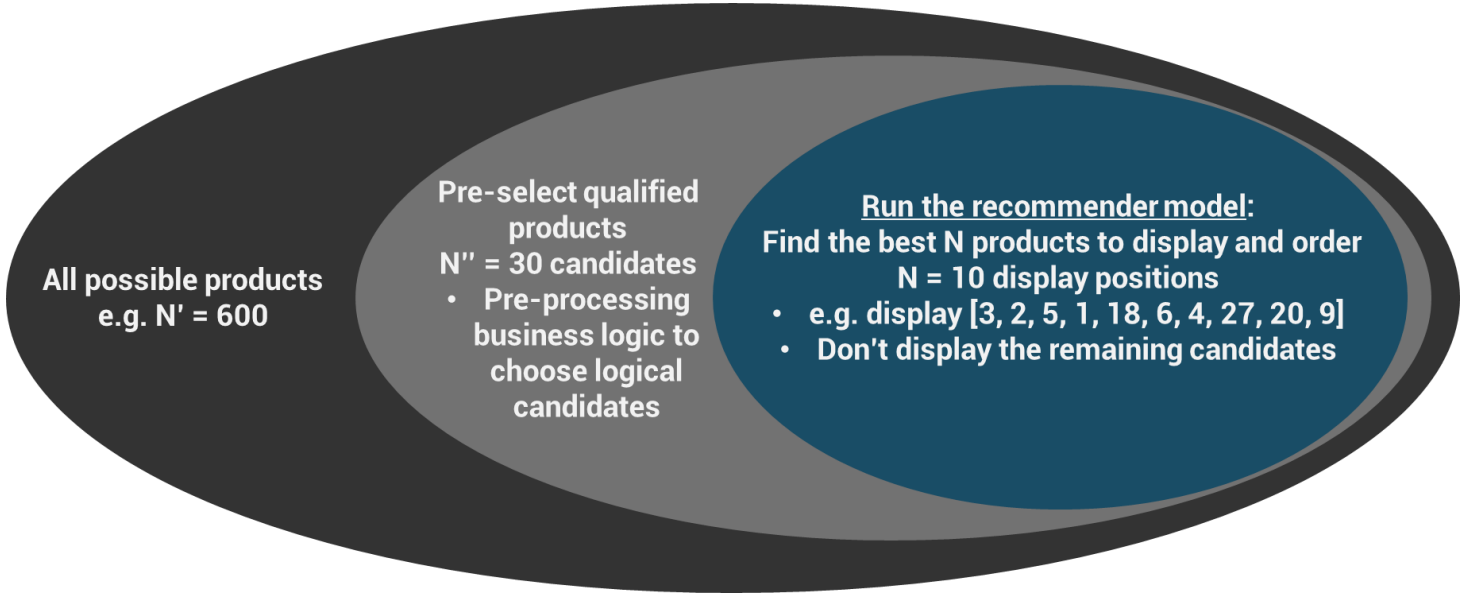


The objective of the recommender model:

Given a list of candidate products applicable to a web page, build a 'self-learning' model which can select the best ones to display (in the correct order) to a given customer segment

# Selecting and displaying products

Our models are currently being used to answer these business questions:

Which specific products are best displayed for this customer segment?
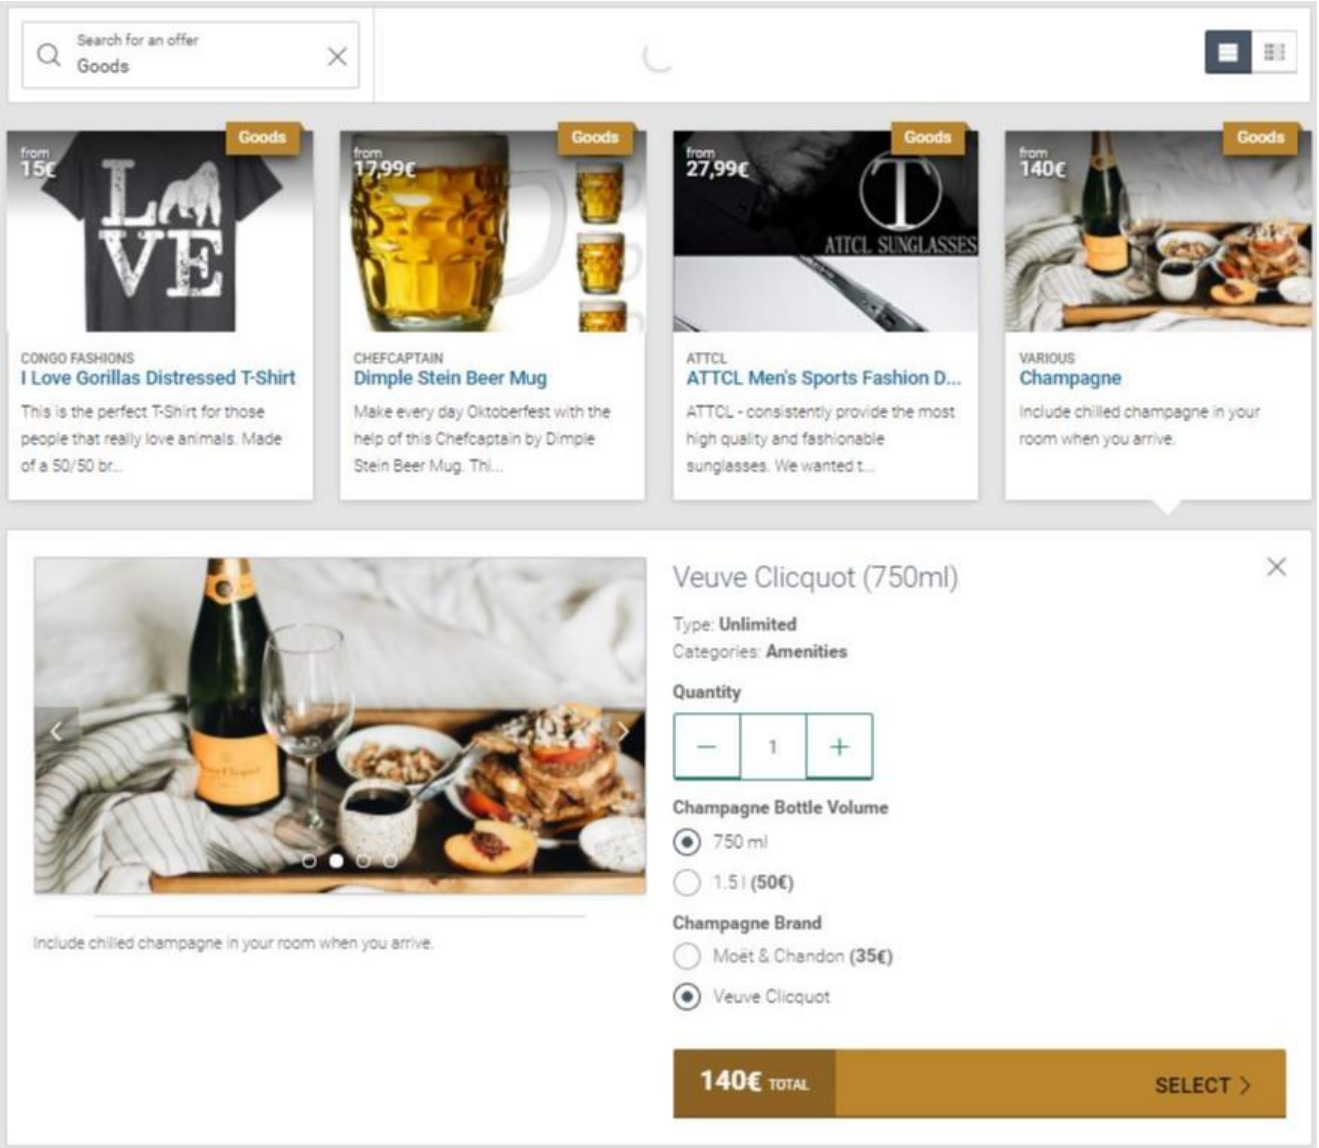
What is the best display order of those products?