

Data Science: Key Concepts

In this chapter we will also look at the five disruptions that are caused in the market place by data science. Once the context and its importance is understood it's easy to simplify and demonstrate what data science actually is. We will also study traditional architecture versus Data science and understand the importance of Signal detection, which we shall study as chapter 2 and the machine learning techniques that help with this signal detection is studied from chapter 8 onwards, although we have covered few machine learning concepts in this chapter. This chapter shall also discuss solution architecture and the three critical components that are required for any solution.

FIVE DISRUPTIVE PRODUCTS

The five quick disruptive products launched in the market place will be discussed now:

1. A very simple Japanese App
2. Healthcare App
3. Coursera
4. Sensory device in Agriculture Sector
5. Autonomous Car

THE JAPANESE APP

The first one is a very simple Japanese app, which essentially helps two people to discover each other. Essentially, what the App does is, for every individual a set of questions has to be answered. When these questions are answered it gives a characteristics score that tells if the person likes music, books, viewpoints on philosophy, religion etc. Whatever the parameters are, the questions have to be answered and each person gets a score attached to each question answered.

The other score that is attached to this device is the location. If a device is carried while walking on the street it will tell how many people with similar scores are around you within a 1 km radius. This app will enable strangers to look up at one another and have coffee, chat or get to know one another better. Using similarity score and location they are able to discover one another.

Disruption: An app that leveraged and capitalized on new social norms of today's casual meetups. Revolutionized the way people find others with similar taste/interests. Usage of data to find patterns and clusters from humongous set of entries and present to the users in a meaningful way, which is 'right match' in this case. Turning Data to Insights.



FIGURE 1.1 Japanese dating app

THE HEALTHCARE APP

The second one is in the healthcare space. In this healthcare app a heart implant is able to communicate information such as rate of heartbeat, condition of heart in real time with your mobile phone. The mobile app also communicates remotely to the doctor.

Disruption: Reduction in visits to the clinic, reduction in non-medical costs. Continuous monitoring of organ health vs. one time data captured during the physician visit. Presents an opportunity to track patterns and higher chance of identifying an anomaly and hence act early/on time.



FIGURE 1.2 Heart implants

COURSERA

The third disruptive product is Coursera, an online educational platform where one can learn various kinds of courses for free. There are a lot of educational videos and tutorials online. When students watch these videos it is possible to pinpoint those places in the video when students pause or stop. Those jump and exit points are noted and this enables to figure out how to re-orchestrate the content, to make the content more engaging.

Disruption: While MOOCS have expanded the access to education to learners by overcoming lack of infrastructure/resources, COURSERA aimed

coursera Courses Specializations Institutions

Take the world's best courses,
online, for free.

What would you like to learn about?

Join 8,423,512 Courserians.
Learn from 688 courses, from our 110 partners.
[How it works »](#)

FIGURE 1.3 MOOC

to continuously improve the quality of the content delivered by collecting data on focus/topics of interest from thousands of students from across the world. By redesigning UX, and fine tuning content COURSERA disrupted the way online education was delivered by its predecessors like Khanacademy, MIT OCW, etc.

SENSORY DEVICE IN AGRICULTURE SECTOR

Fourth, disruptive product is in the Agriculture sector. Netherlands agriculture is a big part of their economy. They make the worlds best cheese and butter. One of the problems farmers face there is understanding the health of cows, which are carrying. Therefore now they have attached a sensory device to the cow's ears, through which farmers can remotely (communicated via a satellite), monitor their cow's health.

Disruption: Livestock farming techniques and the sensors help with cattle health monitoring and action can be taken immediately if the cattle are unwell. This helps within time detection of disease and helps prevention of spread of disease to the other cows through prediction.



FIGURE 1.4 Sensored cows in Netherland

AUTONOMOUS CAR

Lastly, the autonomous car, an autonomous car is special in that the car moves without a driver. This device tracks and scans the surroundings of the car at high speeds. It has the intelligence to process all kinds of real-time information and communicates it back to the steering wheel.

Disruption: Processing data from images and supplementary sensors, self-driving cars create a virtual world through which they navigate. By reducing the reaction time by millions of folds than human level, they aim to eliminate human error driven accidents and traffic congestions. Significant improvement in time and fuel efficiency whilst saving lives.



FIGURE 1.5 Googles autonomous car

A look at all the five uses shows one thing that is common to all of these and that is a data product which is working behind the scenes, very silently humming. To create a data product a data science process is needed, which will unlearn patterns from that data and create a bigger product. So in the five examples that happen in our everyday like how our health gets taken care of, how we learn, how we fall in love, how we farm and how we drive, all of these are touched increasingly by data products. Data science needs to be an integral part of any organization you consider, else there is a very high probability that you will lose the market place.

One of the biggest secrets of winners is that they are able to see patterns faster. So a core team, which uses data science techniques to process all the structured, unstructured data and looks at patterns around it and acts on it in real time is what most companies are aiming at today.

DATA SCIENCE Vs TRADITIONAL METHODS

It's similar to an iceberg floating on water. Most organizations just see the tip of the iceberg. For example they just know how much sales is happening.

They fail to realize what is driving sales. If there is a change in the promotions by 5% what is the expected growth in sales? There are lots of unknown questions for which answers are required.

Most organizations have tons of data on sales, finance aspects; call centre data and reports, which are typically delivered on Business Objects, Cognos, and Microsoft Analysis Services. These reports quickly answer few important basic questions such as which call centre agent has the best all round time. What happens in Data science is inserting a process called analytical modeling process where there are specific techniques such as segmentation, scoring models, text-mining models, which will process the data and give a different lens. This will enable one to see patterns in the data.

DIFFERENCES IN ARCHITECTURE

Here is a detailed architecture of traditional companies versus the new age companies. Both of them have a Data Repository and a Dashboard but where they are different is in the four layers. There is Machine Learning Process (Text Mining, Collaborative filtering) in-between the data repository and Dashboards, which will change the game. They detect what is called a signal. A Signal is nothing but a pattern, so once the pattern is detected via an action, they keep a close watch on that action. This is a simplified view of the Data science architecture.

What are 4 core differences between DataScience & Dashboards?

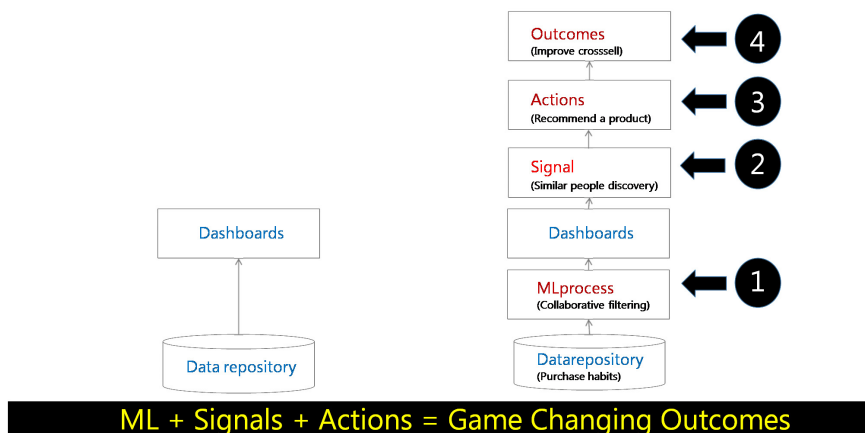


FIGURE 1.6 4 core differences between data science and dashboards

DEMYSTIFYING MACHINE LEARNING

The goal of Data scientist is to use data to discover signals that cause changes and which ultimately have an impact on the revenue of the firm. Even for a data scientist, it is humanely impossible to analyze big data. But with the aid of a computer, it can be easily done. Yet, a computer can only compute what has been programmed into it. So how do data scientists cope with this scenario, where analysis of the data will require the computer to pick up the ‘trends’ on its own? This is where machine learning comes in.

Machine Learning is a remarkable application of artificial intelligence that enables computing systems to perform tasks through a process of “self-learning” without their being specifically programmed for the same. As data scientists cannot pinpoint exactly what sorts of patterns, the computer should recognize, this application of “machine learning comes in extremely handy. Thus, machine learning facilitates the computer to automatically adapt to new patterns and signals in data, while “learning” or recognizing previous trends and data computations. When Google’s search bar uses “auto-complete” before you type in your query, it is an example of machine learning, as the Google server has learnt to give you ‘predictions’ of what you might want to search based on your previous search history.

We will now familiarize with five techniques

TECHNIQUE 1: SEGMENTATION

This process involves breaking data into various chunks based on shared characteristics. The analyst then picks the clusters through an iterative process looking for uniqueness between segments. We could segment based on demographic, need based, behavior based etc. The statistical techniques that we use for segmentation are K Means, Hierarchical clustering and Discriminant analysis, as shown in figure 1.7.

Some business questions that are answered by segmentation are:

- What are the behavioral personas about customer, which lie buried in my raw customer transactions in the database? This is explained in Figure 1.8
- Which specific customer behavior discriminates a high value segment from low value segment? This is explained in Figure 1.9
- How do customer behavior segments migrate across time and what does it reveal to us? This is explained in Figure 1.10 and 1.11



- **Customer Context**
 - A large owner of fleets in US
 - Each truck driver given a fuel card
 - Driver info + Mileage + Refuelling behaviour + Location
- **Customer Challenge**
 - Aligning Service Models to Customer Segments
 - Drive Growth & Ability to Cross-sell & Up-sell
- **Data Science Technique**
 - K means clustering
 - Analysed over 120,000,000 Customer Records & Profiles
 - Analysed over 110,000,000 Million Customer Service Rep Comments

FIGURE 1.7 A Real life customer segmentation case study

Fleet related master data

1. Fleet id
2. SIC Code
3. No of trucks in fleet
4. No of drivers/cards

Fleet spend data

1. Avg_Gallons_Per_month
2. Avg_Spend_on_non_fuel
3. Avg_Transaction_Per_Month
4. Total_Active_Cards
5. MOM(3 months) growth(gallons)
6. Avg_Credit_utilization (3 months)

Current product holdings flag

1. Has_OPIS_Suite_of_Reports_flag
2. Has_EFPS_Discount
3. Has_Smart
4. Has_Rewards
5. Has_Screen_Now_Report
6. Has_Volume_or_Service_Discount
7. Has_Exception_Reports

Touch point data

1. Avg No of inbound calls per month
2. Recency of last call
3. Total no of phone calls per year

FIGURE 1.8 Behavioral components considered for fleet card segmentation

1	Av. Gallons ppm	Av. No of Transactions ppm	Av. Ancillary Revenue	Av. No. of Retention Calls	Av. Late Fees	Av. Activation Rate	Av. MOM Growth
Population	106.4	5.5	248.8	0.7	324.1	0.7	1.03
Stable Underdogs	62.5	3.6	83.4	0.3	87.7	0.5	1.05
Miniature Laggards	77.3	4.6	90.0	0.2	118.3	0.9	1.03
Cash Cows	179.9	8.0	1098.6	0.6	2965.3	0.7	0.96
Dark Horses	122.3	6.1	1196.2	0.4	534.6	0.7	0.93
Sulking Mediocres	101.5	5.2	63.2	4.6	279.2	0.6	1.07
Front-runners	276.4	12.4	163.8	0.7	371.0	0.7	1.04

Note:

	Unstable Behavior
	Average
	Stable Behavior

FIGURE 1.9 Dimensions of fleet behavior measured and segmented

Segment-3: Cash Cows...Segment Profile

- Definition: The large size fleets, that are mostly medium tenure customers having very high spends but also having high late fees incidences
- Constitutes 5% of total fleets and contributes 22% of total spend.




FIGURE 1.10 Cash cow - segment profile

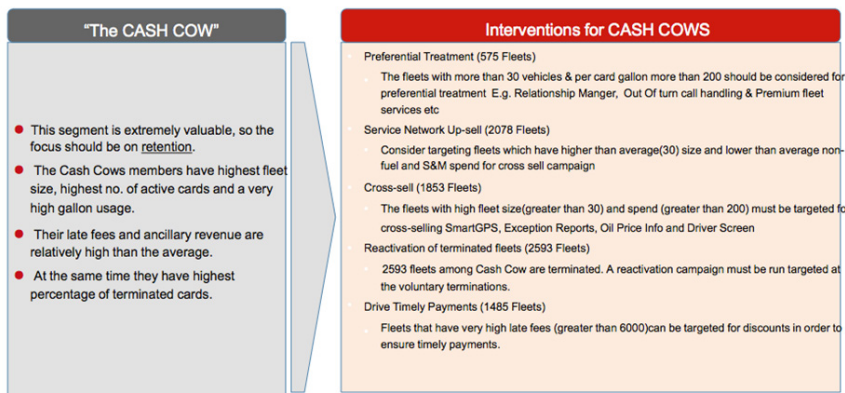


FIGURE 1.11 Cash cow – behavior portrait and target action

Segmenting in BANKING Industry

In order to give the right offer and product to the right customer and to do it the efficient way you will need to use a segmentation method. In banking we could classify and segment the customers into 5 clusters and their line of credit, pricing and campaign intervention for each segment can be studied as seen in the graph 1.12

Clustering

It is considered the most important unsupervised learning problem. Cluster analysis is in simple language dividing data into different clusters or groups.



Key cluster observations

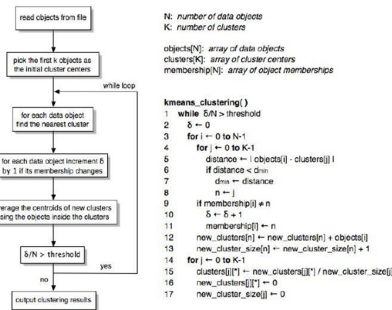
- Cluster Observation 1: Low balance, Lowrisk, Reached credit limit often
•Possible treatment strategy: Extend Line of credit and possibly charge fixed fee depending on # of times they reach credit limit
- Cluster observation-2: Low balance, moderate risk, reach credit limit often
•Possible treatment strategy: Possibly charge fixed fee depending on # of times they reach credit limit
- Cluster observation-3: High balance, moderate risk, Do not reach credit limit often
•Possibly run a focused outbound campaign to sell short term fixed deposit
- Cluster observation -4&5: Moderate balance, High risk, Moderate usage
•Since risk is high, interest rates and Pricing strategy

5 segments and LOC, Pricing, Campaign interventions for each customer segment

FIGURE 1.12 Segmentation in banking industry

The greater the similarity within a group the better is the cluster. The greater the dissimilarity between groups the cluster is more distinct. One technique of clustering is the k means technique. This technique is used to separate data into the best-suited group based on information the algorithm already has. Once data is separated one has to specify the number of cluster that will be created to be able to produce effective data mining results. Each cluster had a centre point called the centroid, which each observation is assigned to. Associating every observation with the nearest mean creates K clusters.

Then one has to calculate the centroid mean for each cluster. This becomes the new mean and the above two steps are repeated till convergence has been reached.



- K means algorithm
- Specify K the number of clusters to create
- Choose K points at Random as Cluster centroids
- Assign each observation to the cluster centroid it is closest to
- Calculate centroid mean for each cluster
- Use it as the new centroid of the cluster
- Iterate till cluster centre does not change

FIGURE 1.13 The mathematic behind clustering

TECHNIQUE 2: UNSTRUCTURED TEXT MINING

The second technique is unstructured text mining. Here we use data to discover signals and process changes that create an impact. So let's take a store manager as an example. A store manager gets a lot of feedback and there is a lot of unstructured data that comes in. It's very important to process this into a structure. Text mining can process all this feedback and give a glimpse of what is called the sentiment analysis as how many people like the store and how many don't like the store.

Another example of text mining is in the health care domain. For example, a patient goes to a doctor and gets admitted to a hospital. He comes in contact with a doctor, lab technicians and the nurse. The Doctor makes a record of the state of his health condition, the Lab technician tests his blood and writes down some inferences and the nurse regularly checks his vital parameters and makes notes of the health condition of the patients. Now all this data can be run through a text mining activity and triangulate the state of the patient by mashing up all the three datasets, the doctors dataset, the lab technicians dataset and the nurse's dataset.

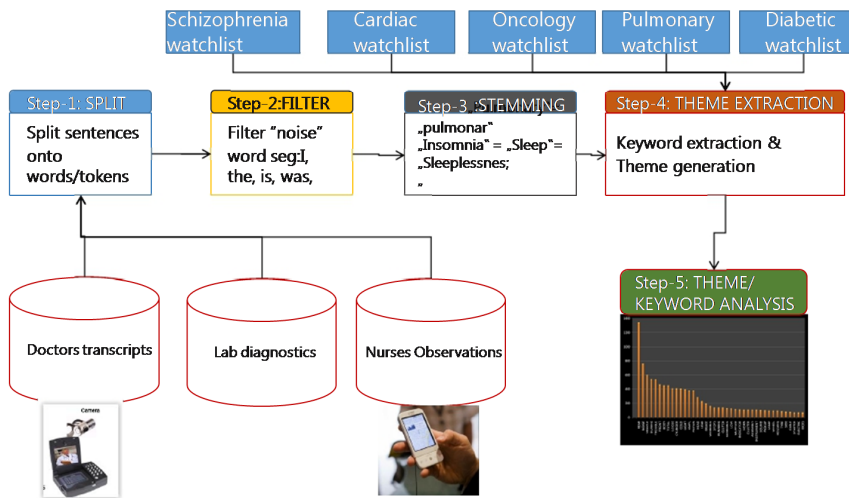


FIGURE 1.14 Real world unstructured text mining in the health care

Another example of text mining is insurance domain. Auto insurance companies deal with large number of claims every day from collision damage, fire and theft damage, accidental damage. Large amount of time and money is spent in identifying fraudulent claims. One of the datasets used is historical claims, coverage and settlements. A text mining solution

layer is added above the structured data that quickly reads and understands the claim details to highlight missing facts, inconsistencies and changed stories to identify with a probability that the claim is fraudulent.

TECHNIQUE 3: USE OF SCORING FOR SIGNAL PROCESSING

The third technique is scoring models. Scoring refers to what is the best action that a person is likely to take? For example most of us have shopped at Flip kart at least once. Flip kart tries to figure out what is the next best action that a customer is likely to perform; will one be redeeming a coupon or is the customer just browsing and adding items to the wish list? Studying the past history and trying to characterize the future behavior is a scoring model. It is a binary result, its either an action or not, in other words 1 or 0.




FIGURE 1.15 Scoring models

TECHNIQUE 4: FORECASTING


Fourth technique is Forecasting. When trying to predict the sales for the next three months, if the outcome is a number then this is Forecasting as seen in Figure 1.16.

TECHNIQUE 5: RECOMMENDERS

The fifth technique is Recommenders. Looking at recommenders, if a person likes pizza and salad and another person likes pizza, salad and coke. Based on the actual purchase behavior one will club these two people as similar to each other. As both people have similar purchase behavior and only one has purchased coca cola the other person will be recommend or offered coca cola.

**FIGURE 1.16 Forecasting**

Recommenders are one aspect of data analytics that most of us in some form or the other would have come in contact in our day-to-day activities. When you buy items from a grocery store you could generally get coupons to use for next purchase. When you visit a salon, you could be offered a combo deal with manicure / pedicure along with your hair cut rather than just haircut. These generally happen because the company knows that people who buy a set of products, for e.g., eggs and milk, could buy bread and cheese next and hence recommendation happens via coupons. In the past new comers who have been to the salon have generally tried manicure along with a haircut so when you for go the first time the company markets both services to you.

**FIGURE 1.17 Recommenders**

Snapshot of Machine Learning Techniques

As in the above topic we have just seen few machine learning technique below is a graph that displays a snapshot of six techniques and a brief idea of what falls under each of the techniques. We will study Machine learning as a whole topic in chapter 8.

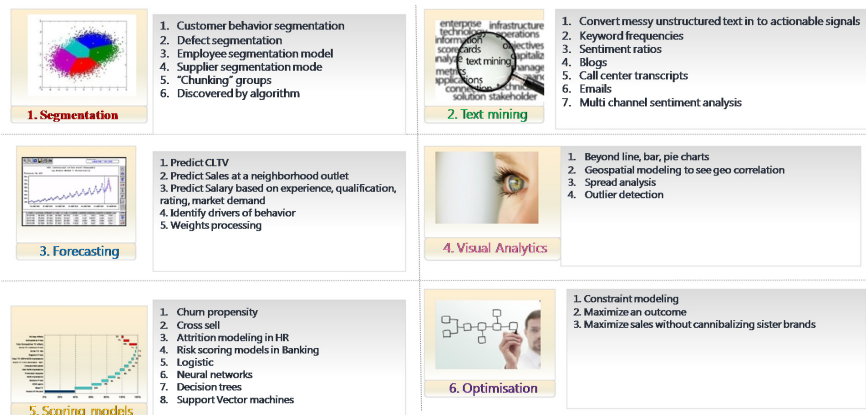


FIGURE 1.18 Machine learning reference architecture

REFERENCE ARCHITECTURE

A look at the reference architecture for any of these projects has three layers. Machine Learning Reference Architecture is an algorithm that can process raw data to provide a big picture that combines all the major and minor aspects of the data being analyzed. Real time data analysis is carried out using an optimization process,

1. The store layer that captures and stores data: where one can store the data in Hadoop, Hive, Hana or any other database. More important than storing is what is done with the data, how to extract signals from the data and that's where data science comes in.
2. The Sense Layer: in this layer a text-mining model or a scoring model is used to detect a pattern. Then it mines the collected data from historical trends and patterns that act as reference points, this pattern is then monetized.
3. The Respond Layer: compare previous trends with the latest data collected, to predict an outcome and to recommend the next decision to be taken. This is where the analysis of the data is presented in relatable

terms, and the patterns are contextualized in terms of the business/company's needs. Analytic workflows are computed and insights are generated.

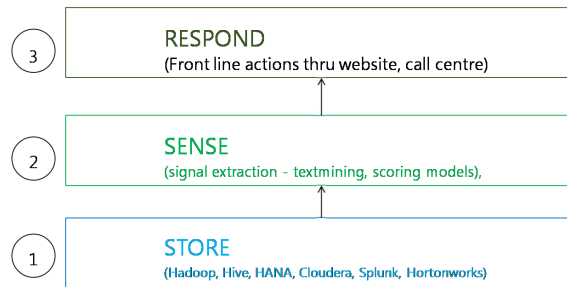


FIGURE 1.19 Machine learning reference architecture

Architecture of a Data Product

At the broadest level, First thing to answer is the Robust Data Ingestion. Every click must be ingested. Lets say on Amazon or Netflix, millions of clicks happen. Each click must be processed and scaled and stored in a robust way. In the year 2000 or a little later Netflix started to find a different way forward to understand the customer preferences. Rather than predicting stars, Netflix made an effort to predict genres.

Netflix is taking movie metadata and categorizing using the micro categorization strategy with over 76,897 micro genes, enabling fine-grained recommendations of what type of movies a person will like. For example My Netflix categories show movies under these types:

- Comedy, Drama
- Drama, Romance
- Romance, Comedy
- Musical, Fantasy, Comedy

Lets say one has varied categories of his interest like Action, Adventure, and Cartoon. When Netflix makes a recommendation it bases it on the kind of movie a customer wants to watch rather than the rating its customer assign to a particular movie.

Second thing is to run an analytical process wherein one runs text mining, and tries to cluster the customer to understand their behavior. Based on all of these analyses, action is taken.

Where do these genes come from? This is something that Madrigal went on to do an extensive study and using textual analysis; he made findings

on keywords usage frequencies in a visualization format producing some accurate and some not expected findings. At the foundation of each micro gene is a broader genre, like romance, comedy or drama. From here Netflix then has other advanced filters like location, time period, actors to get to the precise category. Since Netflix uses algorithms to cater to its customer base, a human analyzes the films and then its put through a Netflix process, which brings certain patterns. Example micro genes about friendship and marriage were among the most common. He also discovered that Drama, classic were some of Netflix's widely used adjectives for movies.

HANDS ON SEGMENTATION

You will learn the path to download R in chapter 2, here you can have a quick view of some commands in R specific to clustering technique and its visualization.

1. To segment the dataset we use the “kmeans” command to find clusters. To view the cluster statistics we input the “fit” command.

Hands on Segmentation... Using K means to find clusters

```

> fit <- kmeans(retail_data_clustering, 5)
> show(fit)
K-means clustering with 5 clusters of sizes 61, 35, 78, 74, 52

Cluster means:
  AvgPurchValue  AvItemPurchased  ShopFrequency  Spendtilldate
1      38170.67           6.803279      3.622951      50209.34
2      48370.06          12.771429      8.457143      140760.31
3      49059.06          11.846154      5.487179      96592.00
4      17516.01           9.756757      5.621622      81245.81
5      16695.88          12.326923      8.173077      133726.56

Clustering vector:
 [1] 4 4 4 2 3 2 3 3 2 3 4 3 3 5 4 4 4 3 3 1 3 2 3 5 4 4 5 3 2 5 5 2 3 3 1 2 5 2
 [38] 5 3 1 4 2 5 4 3 4 1 5 3 2 2 4 2 4 4 4 1 4 5 3 3 4 1 1 4 1 3 1 2 1 3 5 4 5
 [75] 5 3 1 1 3 4 3 4 3 2 3 5 1 3 4 5 2 5 2 1 2 1 1 1 1 1 5 1 4 2 3 4 4 3 4 4 3
 [112] 1 5 5 4 1 5 3 1 3 3 4 3 1 4 4 1 5 1 2 1 3 4 5 4 3 5 5 4 4 5 3 1 1 1 1 4 5

```

Execute clustering algorithm

Shows cluster statistics


Shows membership of each segment

Is it above or below population average? How does this help characterise the segment?

2. To visualize the data we use the “ggplot” command. On the X axis we have plotted average purchase value and on the Y axis we have plotted average items purchased for two cities Bangalore and Chennai.


ggplot() – How to draw a quick scatter plot? Visual relationship

```
> ggplot(retail_data_clustering, aes(x=retail_data_clustering$AvgPurchValue,
+ y=retail_data_clustering$AvItemPurchased,color = retail_data_clustering$City ))+
+ geom_point() +
+ geom_smooth()
```




- Another way of visualizing is in 3D format as shown in the figure below using the “scatterplot3d” command.

```
> scatterplot3d(retail_data_clustering$AvgPurchValue,retail_data_
+ clustering$AvItemPurchased,retail_data_clustering$Spendtilldate)
```



4. We could draw a box plot to analyze the spread by inputting the “boxplot” command for the retail data.

```
> boxplot(retail_data_clustering$AvgPurchValue,  
horizontal = TRUE, xlab = "Average Purchase Value")
```



SUMMARY

Any organization today is drowning in data. They have tones of data with them. Its clearly evident in the job market that the role of a data scientist is scare and a look at the job portals shows how companies such as KPMG, Pfizer etc., are all recruiting data scientists. Data is an asset that can add value to their organizations, and can help them predict future market demands. Using this data is a must to improve decision-making, and it can only be done through data analysis. Traditional methods of analyzing data can no longer be used because of the huge amount of data flowing in, making it unavoidable to use Data Science techniques, and specifically Machine Learning.

While data mining has been used earlier to discover previously hidden information and anomalies, Machine learning goes one step further by using the stored information collected previously to predict future trends. Undoubtedly, Data Science is indispensable to businesses that deal in share markets, advertising, political campaigning, voter behavior and advanced medicine or any other domain. With more and more efficient algorithms being developed, it is possible to filter the necessary information in today's data-rich world that can aid in making vital decisions.

As an introduction to data science we have understood why data science is important, and the option to make this as a career as data products surround us. We have also touched five situations to understand this. We have also understood the core difference between data science and a traditional dashboard being, the Machine learning process, which is the critical thing for a data scientist to work on. We have studied five concepts of machine learning to familiarize and also seen few examples from the industry. We have also seen the Machine learning reference architecture.