



Welcome to Business Analytics Quiz Accendo 2018 BRIDGEi2i Analytics Solutions

QUIZ MASTER: TUHIN CHATTOPADHYAY, PH.D.

ORGANIZER: MS. BHARGAVI SHANKAR

10TH JANUARY, 2018

Quiz Format: 5 Rounds

1. Analytics Business
2. Machine Learning
3. Big Data
4. Audio Visual
5. Buzzer

First 4 Rounds:

- ✓ 2 Question for each of the 5 teams
- ✓ 4 Options available

Last Round:

- ✓ 10 questions will be fired at all the teams one after another
- ✓ No choice in this round



Rules for First 4 Rounds

- ▶ 10 marks will be awarded for a correct answer at the first chance.
- ▶ 5 marks will be deducted for each wrong answer.
- ▶ If a team cannot answer a question, it can be passed to the next team.
- ▶ If answered correctly 5 marks will be awarded to the team or else 2 marks will be deducted for a wrong answer.
- ▶ Only one pass is allowed for a question.
- ▶ Answering time per question is only 30 seconds.
- ▶ Discussion within the team is allowed.

Rules for Last/ Buzzer Round

- ▶ A question will be posed to all teams and the team which presses the buzzer first will be given the chance to answer.
- ▶ 10 marks will be awarded for each correct answer.
- ▶ A deduction of 5 marks will be done for wrong as well as no answers.
- ▶ Answering time is limited to 20 seconds per question.
- ▶ If a question is not answered by the first team who pressed the buzzer, the team that pressed the bell next gets to answer.

Round 1: Analytics Business

1. Revolution Analytics was acquired by

A. IBM

B. Salesforce

C. Oracle

D. Microsoft

2. In Gartner's Magic Quadrant for Business Intelligence and Analytics Platforms, Which company tops Niche Players?



A. Citus Data

B. Alteryx

C. H2O

D. Interana

3. Japersoft was acquired by

A. RapidMiner

B. Zoomdata

C. TIBCO Software

D. Databricks

4. Who is the author of the popular book, *“Competing on Analytics: The New Science of Winning”*?

A. Jinho Kim

B. Julia Kirby

C. Thomas Davenport

D. Ian Ayres

5. Baidu, one of the premier AI leaders in the world, has its headquarters at

A. Beijing

B. New York

C. Frankfurt

D. Silicon Valley

6. Hortonworks was formed in June 2011 as an independent company, funded by \$23 million venture capital from two companies. One of them is Benchmark Capital. Which is the other one?

A. Google

B. Yahoo

C. SAP

D. MongoDB

7. Cloudera was founded in 2008 by some of the brightest minds at Silicon Valley's leading companies, including Christophe Bisciglia (Google), Amr Awadallah (Yahoo), Mike Olson (Oracle), and Jeff Hammerbacher. Which company was Jeff from?

A. Facebook

B. AWS

C. Teradata

D. Dell

8. Who is CEO of DataRobot?

A. Jeremy Achin

B. Tom de Godoy

C. Gourab De

D. Jack Collins

9. BI vendor Pentaho was acquired by

A. Hitachi Data Systems (HDS)

B. Amazon Web Services (AWS)

C. 1010data

D. SnapLogic

10. Weka is developed at

A. Australia

B. South Africa

C. New Zealand

D. USA



Round 2: Machine Learning

1. Ridge Regression, Least Absolute Shrinkage and Selection Operator (LASSO), Elastic Net and Least-Angle Regression (LARS) are part of

A. Decision Tree

B. Neural Network

C. Ensemble Models

D. Regularization

2. Exponentiating a beta-coefficient from Cox regression gives you what?

A. Odds ratios

B. Risk ratios

C. Hazard ratios

D. Lift ratios

3. Investigators compared mean cholesterol level between cases with heart disease and controls without heart disease. Which of the following is likely the most appropriate statistical test for this comparison?

A. Two-sample t-test

B. ANOVA

C. Repeated-measures ANOVA

D. Paired Sample t-test

4. DBSCAN Algorithm is associated with

A. Multiplicative Regression

B. Bayesian Statistics

C. Markov Chain

D. Cluster Analysis

5. Fuzzy C-Means Clustering is an instance of

A. Exclusive Clustering

B. Overlapping Clustering

C. Hierarchical Clustering

D. Probabilistic Clustering

6. Hidden Markov Models (HMM) are a special case of

A. Radial Basis Function (RBF)

C. Multi Layer Perceptron (MLP)

B. Dynamic Bayesian Network (DBN)

D. Restricted Boltzmann Machine (RBM)

7. Long short-term memory (LSTM) unit (or block) is a building unit for hidden layers of a/ an

A. Recurrent Neural Network

B. Recursive Neural Network

C. Convolutional Neural Network

D. Auto encoder

8. Invented by Geoff Hinton, this deep learning algorithm is used for dimensionality reduction, classification, regression, collaborative filtering, feature learning and topic modelling.

A. Recurrent Neural Network

B. Deep Belief Network

C. Restricted Boltzmann Machine

D. Convolutional Neural Networks

9. Heteroscedasticity cannot be assessed by

A. Brown–Forsythe test

B. Breusch–Pagan test

C. Mann-Kendall test

D. Goldfeld–Quandt test

10. An ARIMA (p,d,q) model has three parts:
The autoregressive order p
The order of integration (or differencing) d
The moving average order q

Which combination of (p,d,q) represents White Noise Model?

A. ARIMA(0,0,0) model

B. ARIMA(0,1,0) model

C. ARIMA(0,1,1) model

D. ARIMA(0,1,2) model

Round 3: Big Data



1. Splunk can be best associated with

A. Financial Analytics

B. Log Analytics

C. Social Analytics

D. Web/ Mobile/
Commerce Analytics

2. Apache Kafka is used for

A. Service Programming

B. Machine Learning

C. Metadata Management

D. Data Ingestion

3. Facebook Corona, The next version of Map-Reduce from Facebook is an instance of

A. Service Programming

B. Distributed Programming

C. Distributed Filesystem

D. Data Ingestion

4. Apache Ambari is used for

A. System Deployment

B. Distributed Programming

C. Distributed Filesystem

D. Data Ingestion

5. HIPI is a library for Hadoop's MapReduce framework that provides an API for performing

A. Signal Processing

B. Sound Processing

C. Stream Processing

D. Image Processing

6. In Twitter, Hadoop handles batch processing, Storm handles stream processing, and the hybrid system is called

A. Hummingbird

B. Summingbird

C. Mummingbird

D. Nummingbird

7. Hortonworks' HOYA stands for

A. HBase On YARN

B. Hadoop On Y Axis

C. Hierarchical OY
Algorithm

D. Helix On Y Axis

8. Object Oriented Data Technology (OODT) was originally developed at

A. Dell Laboratory

B. NASA Jet Propulsion
Laboratory

C. Intel Laboratory

D. GE Laboratory

9. The SWIM benchmark (Statistical Workload Injector for MapReduce), is a benchmark representing a real-world big data workload developed by

A. Stanford University

B. Princeton University

C. Yale University

D. University of California

10. Oozie is meant for

A. Log Collector

B. Data Exchange

C. Scripting

D. Work flow scheduling

Round 4: Audio-visual

ALPHAGO ZERO CHEAT SHEET

The training pipeline for AlphaGo Zero consists of three stages, executed in parallel

SELF PLAY

Create a 'training set'

The best current player plays 25,000 games against itself
See MCTS section to understand how AlphaGo Zero selects each move.

At each move, the following information is stored



The game state
(see 'What is a Game State' section)



The search probabilities
(from the MCTS)



The winner
(+1 if this player won, -1 if the player lost - added once the game has finished)

RETRAIN NETWORK

Optimise the network weights

A TRAINING LOOP

Sample a mini-batch of 2048 positions from the last 500,000 games

Retrain the current neural network on these positions

The game states are the input (see 'Deep Neural Network Architecture')

Loss Function

Compares predictions from the neural network with the search probabilities and actual winner

$$\text{PREDICTIONS} \quad \begin{matrix} P \\ V \end{matrix} \quad \begin{matrix} \text{Cross-entropy} \\ + \\ \text{Mean-squared error} \\ + \\ \text{Regularisation} \end{matrix} \quad \begin{matrix} \pi \\ \text{trophy} \end{matrix} \quad \text{ACTUAL}$$

After every 1,000 training loops, evaluate the network

EVALUATE NETWORK

Test to see if the new network is stronger


Play 400 games between the latest neural network and the current best neural network

Both players use MCTS to select their moves, with their respective neural networks to evaluate leaf nodes

Latest player must win 55% of games to be declared the new best player



WHAT IS A 'GAME STATE'



This stack is the input to the deep neural network

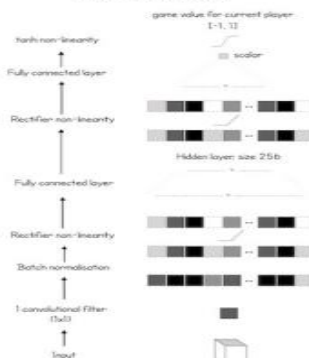
THE DEEP NEURAL NETWORK ARCHITECTURE

How AlphaGo Zero assesses new positions


The network learns 'tabula rasa' (from a blank slate)

At no point is the network trained using human knowledge or expert moves


The value head




A convolutional layer



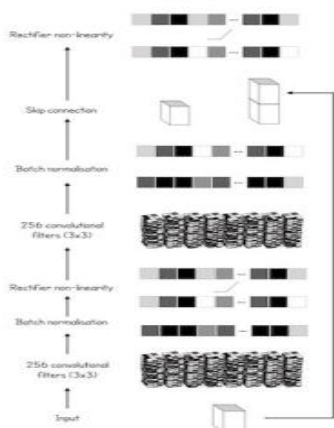
The network



The policy head

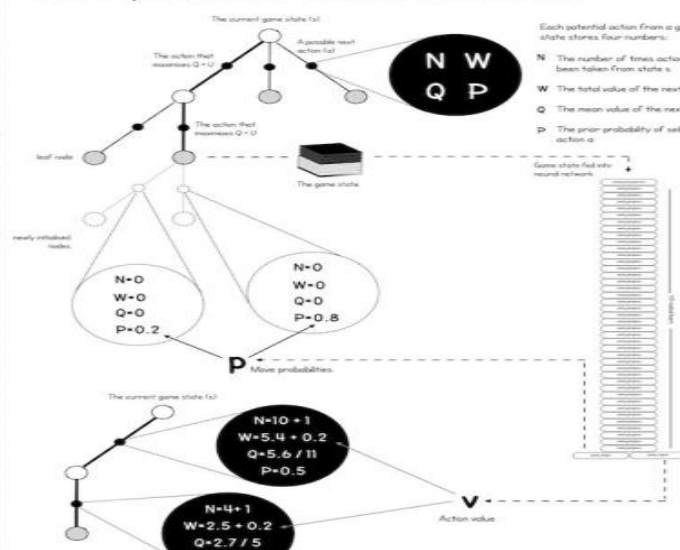


A residual layer



MONTE CARLO TREE SEARCH (MCTS)

How AlphaGo Zero chooses its next move



First, run the following simulation 1,600 times...

Start at the root node of the tree (the current game state)

1. Choose the action that maximises...

$$Q + U$$

The mean value of the next state

A function of P and N that increases if an action hasn't been explored much, relative to the other actions, or if the prior probability of the action is high

Early on in the simulation, U dominates (more exploration), but later Q is more important (less exploration)

2. Continue until a leaf node is reached

The game state of the leaf node is passed into the neural network, which outputs predictions about two things:

P Move probabilities

V Value of the state (for the current player)

The move probabilities p are attached to the new feasible actions from the leaf node

3. Backup previous edges

Each edge that was traversed to get to the leaf node is updated as follows:

$$\begin{aligned} N &\rightarrow N + 1 \\ W &\rightarrow W + v \\ Q &= W / N \end{aligned}$$

Other points

- The sub-tree from the chosen move is retained for calculating subsequent moves
- The rest of the tree is discarded

...then select a move

After 1,600 simulations, the move can either be chosen:

Deterministically (for competitive play)


Choose the action from the current state with greatest N

Stochastically (for exploratory play)

Choose the action from the current state from the distribution

$$\pi \sim N^{\frac{1}{\tau}}$$

where τ is a temperature parameter controlling exploration



1. The diagram is from the paper “Mastering the game of Go without human knowledge” is published in Nature in October, 2017.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel & Demis Hassabis

From which company is the paper published?

A. Amazon

B. DeepMind

C. Microsoft

D. IBM

2. Which science is associated with the movie?

A. Econometrics

B. Bayesian Analytics

C. Sabermetrics

D. Deep Learning



3. What's the area of his doctoral research in University of California, Berkeley in 2003?


A. Reinforcement Learning

B. Financial Analytics

C. Bayesian Statistics

D. Operations Research

4. Name the plot



A. K-M Curve

B. Discriminant Curve

C. Bayesian Curve

D. MCMC Curve

5. Name the Plot

		Case																
		9:USA		5:Japan		8:Ukraine		6:Poland		7:Russia		3:France		2:China		4:Germany		1:Canada
Number of clusters	1	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	2	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	3	X	X	X		X	X	X	X	X	X	X	X	X	X	X	X	X
	4	X	X	X		X		X	X	X	X	X	X	X	X	X	X	X
	5	X	X	X		X		X	X	X	X	X	X	X	X	X	X	X
	6	X	X	X		X		X	X	X	X	X	X	X	X	X	X	X
	7	X		X		X		X	X	X	X	X	X	X	X	X	X	X
	8	X		X		X		X	X	X	X	X	X	X	X	X	X	X


A. Agglomeration Schedule

B. Scree Plot

C. Icicle Plot

D. Dendrogram

6. In a Receiver Operating Characteristic (ROC) curve, the y-axis represents



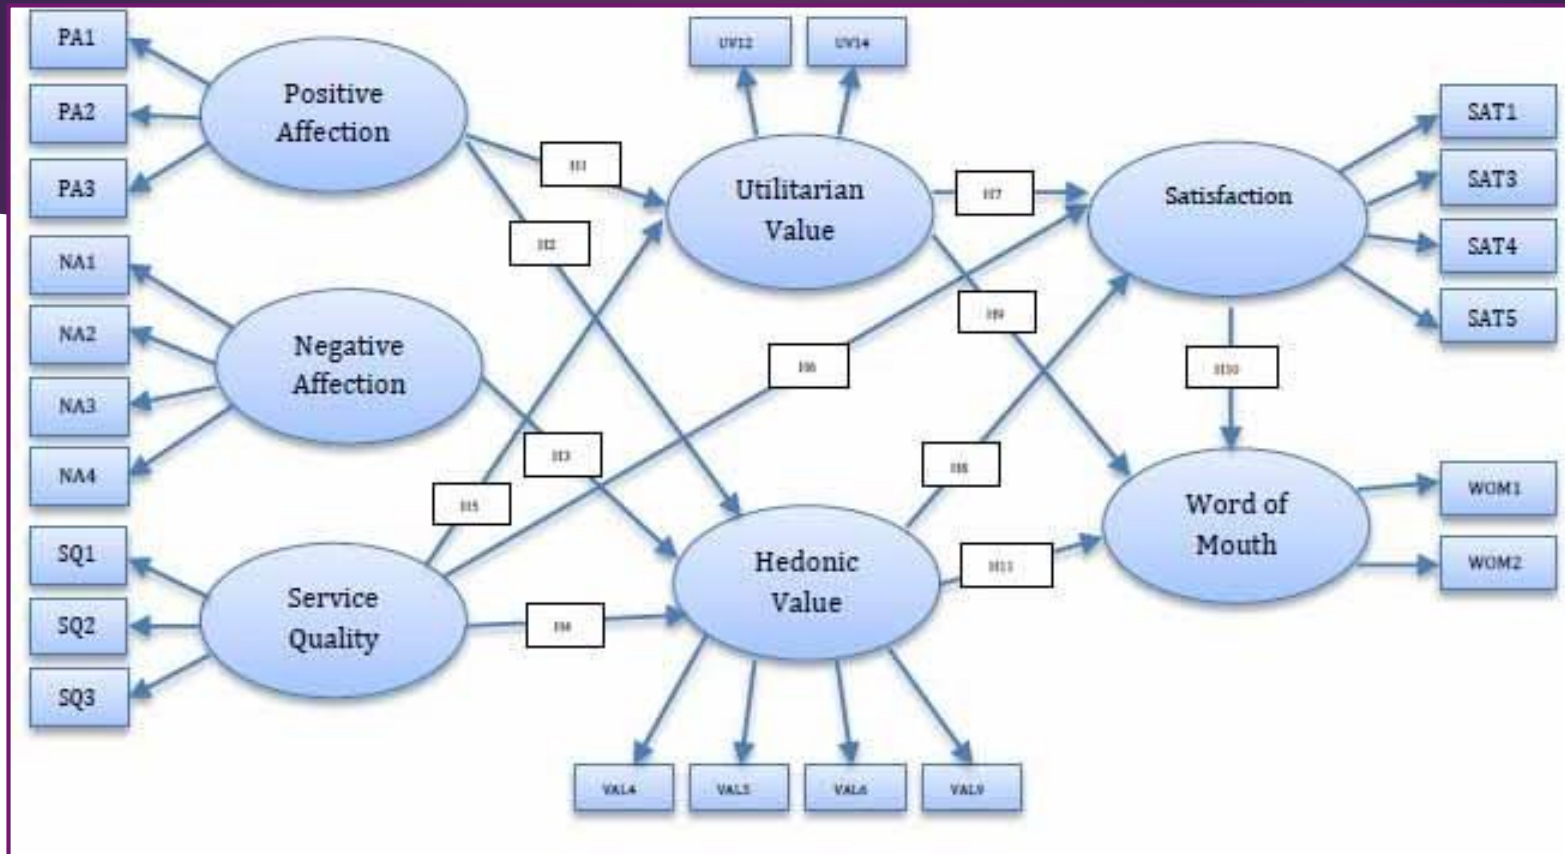
A. Specificity

B. Sensitivity

C. 100 – Specificity

D. 100 – Sensitivity

7. What statistical technique is used in the diagram below?



A. Bayesian Network

B. Markov Chain

C. Structural Equation Modelling

D. Hidden Markov Model

8. Name the Data Scientist in the Video

A. Ros Hansling

B. Shan Roling

C. Ling Roshans

D. Hans Rosling

9. Who is the director of the movie?

A. Steven Spielberg

B. Alex Garland

C. Christopher Nolan

D. James Cameron

10. The protagonist in the movie pioneers

A. Nash Equilibrium

B. Naïve Bayes

C. Network Analysis

D. Queuing Theory

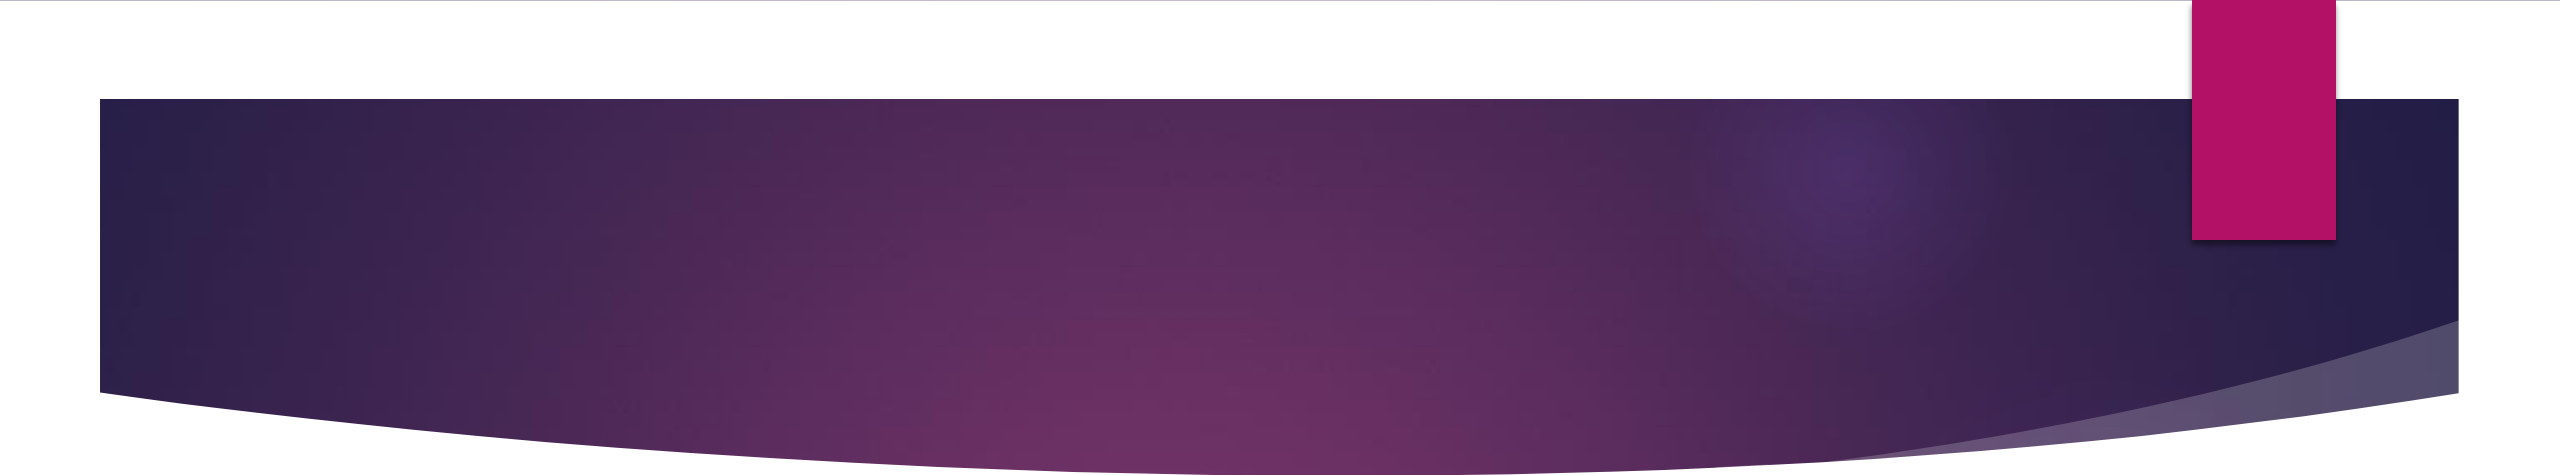
Round 5: Buzzer



1. The impurity (or purity) measure used in building decision tree in CART is



2. What for is Portable Format for Analytics (PFA) used?



3. I am referring to an approach to data collection and analysis in which an automated analytical computation is performed on data at a sensor, network switch or other device instead of waiting for the data to be sent back to a centralized data store.

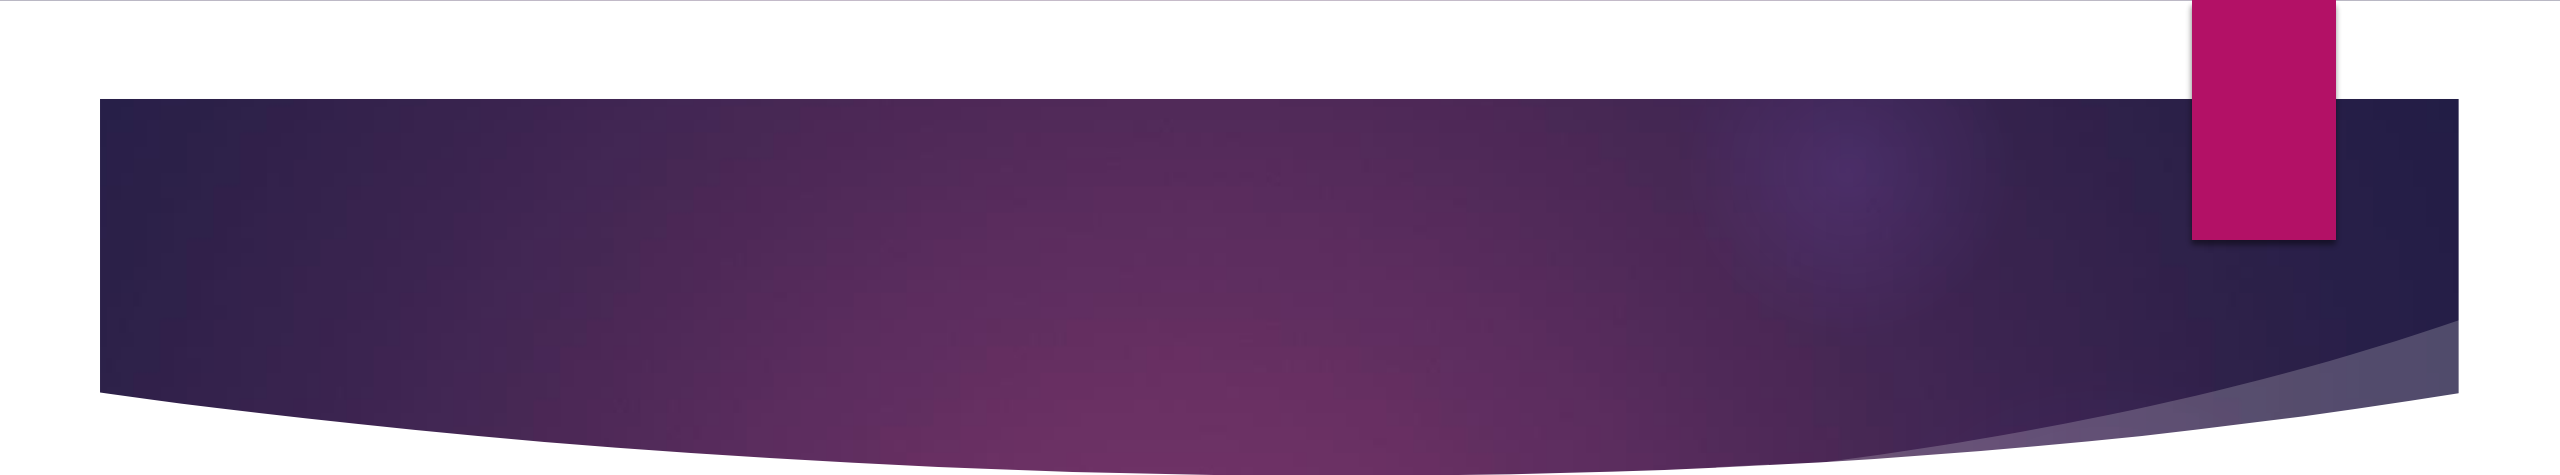
What am I referring to?



4. What does the Levene's test tell you?



5. Mann Whitney is the non parametric alternative
to the_____



6. I am referring to “a special set of scalars associated with a linear system of equations (i.e., a matrix equation) that are sometimes also known as characteristic roots, characteristic values (Hoffman and Kunze 1971), proper values, or latent roots (Marcus and Minc 1988, p. 144).”



What am I referring to?



7. In which type of factor rotation, the axes are not maintained at right angles?

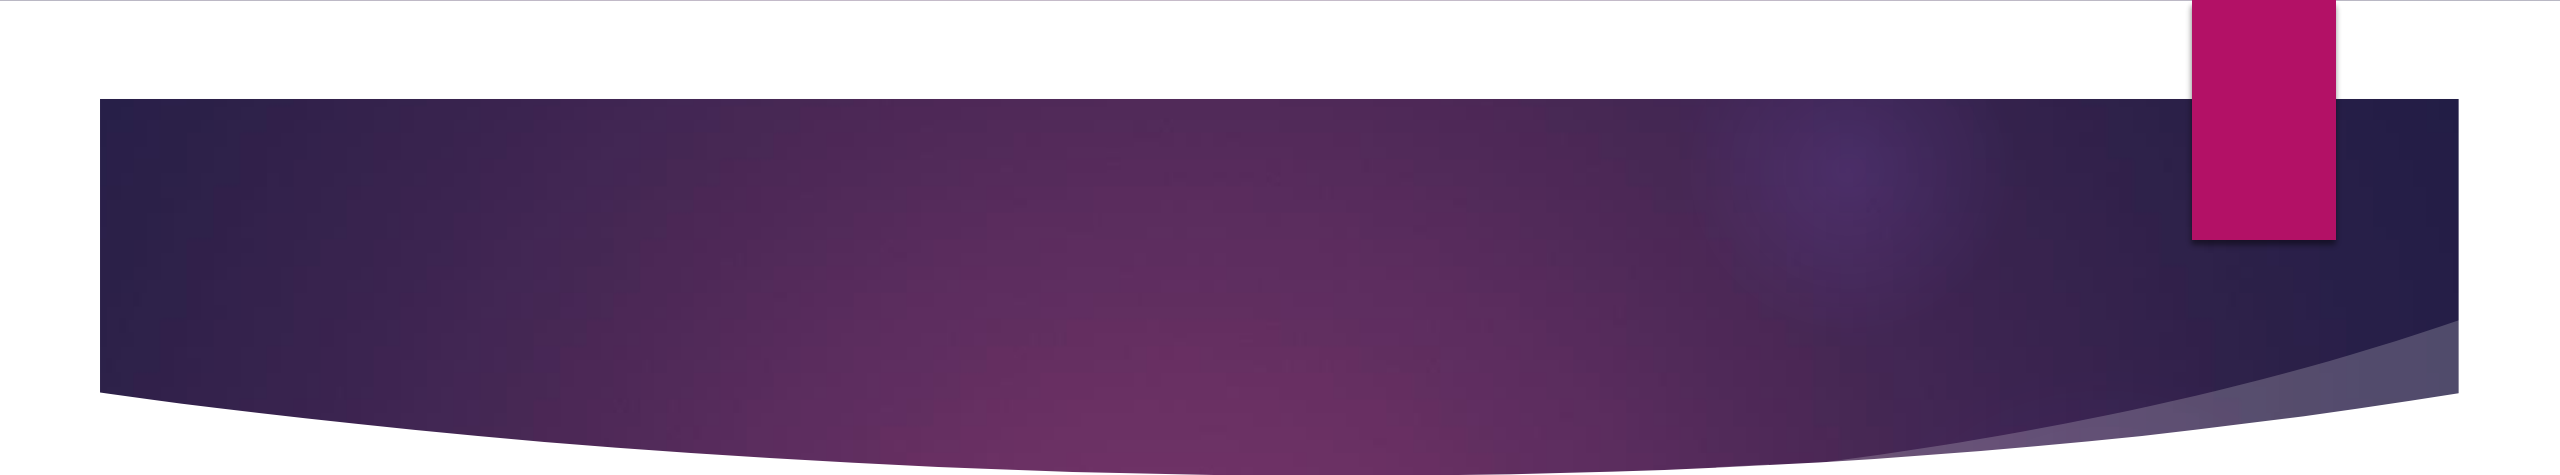
8. I am referring to a variance method (hierarchical, agglomerative) of clustering procedure in which the squared Euclidean distance to the cluster means is minimised.

Which algorithm/ method am I referring to?



9. In Google AdWords Auction, how does Google determine what you pay?

▶ Your Price = $\frac{\text{Your Bid}}{\text{Your Quality Score}} + \0.01



10. Which machine learning technique constructs a hyperplane or set of hyperplanes in a high- or infinite-dimensional space, and uses a kernel function, for classification, regression, or other tasks like outliers detection?