

Come with me if you want to live

Session will start at
19.00 BOT / 13.00 UTC



 A Lunar Cat production...



From the producer of...

SQL Curiosities

Curiosity killed the kitten



THE A.I. WHO TAGGED ME

2020-07-04 19:01 UTC

Message from: ?
"Machine Learning is full of stars.
An Azure Story"



Featuring
André Melancia



How To Be A Human Being - For Beginners



by
André Melancia



Upgrade Your Grey Cells
and use
Azure Synapse Analytics

André Melancia
PDC Conf
2021-09-16



André Melancia's




Murder on
the
Motherboard
Express

2021.
An Azure Bot Odyssey

André Melancia
PDC Summit Pakistan 2021
2021-09-26

The Big Quantum Theory
An Azure Story



André Melancia
Saturday 10th



techforum technical community


Machines Are Learning. You Should Learn ML Too: An Azure Story

Virtual Conference on
Azure Data and AI Platform 2022

Andre' Melancia

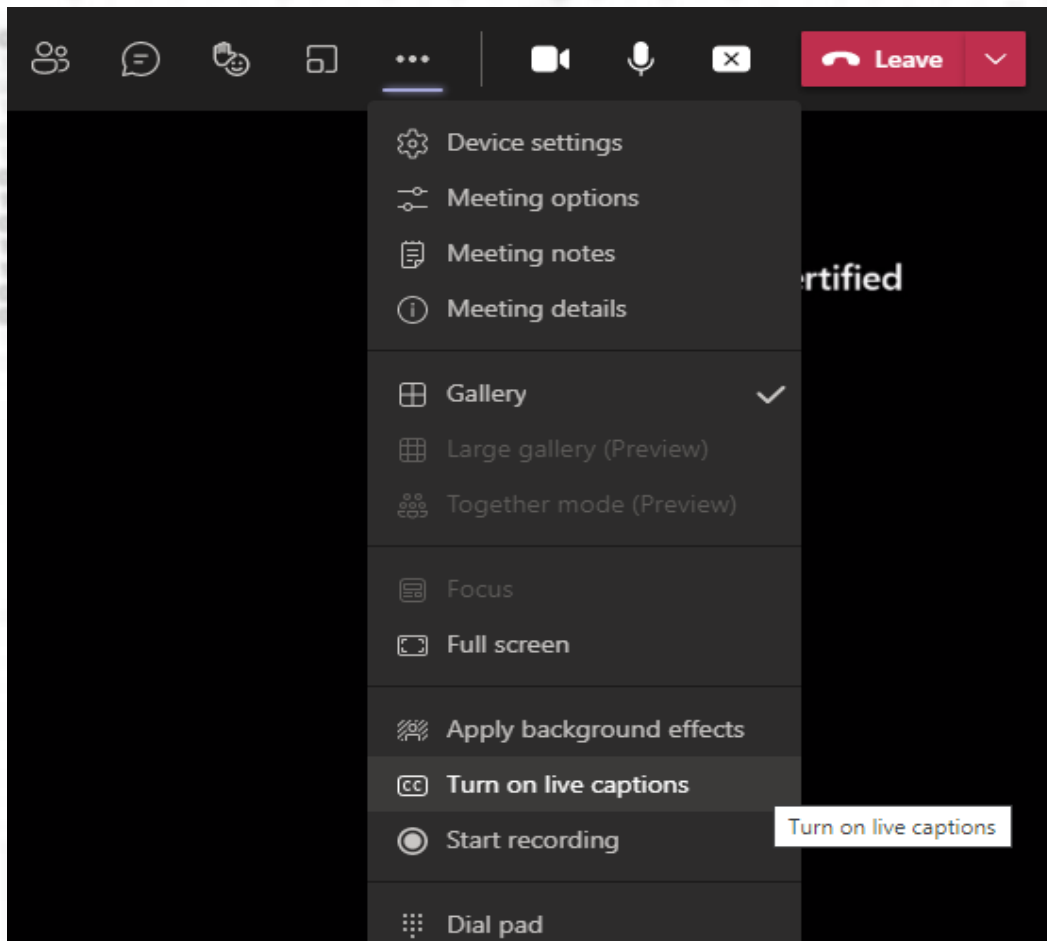
Bangladesh

2022-05-21





AI Subtitles?






Machine Learning \neq Artificial Intelligence


→ Artificial Intelligence

- Rational agent that perceives its environment and takes actions that maximize its chance of success at some goal
 - Solve Problems
 - Make decisions (and do actions) NOT based on previous human programming
-



→ Machine Learning

- Sub-field of Artificial Intelligence
- Study and construction of algorithms that can learn from and make predictions on data.
- Pattern recognition (find patterns in data)
- Predict things





Types of Algorithms

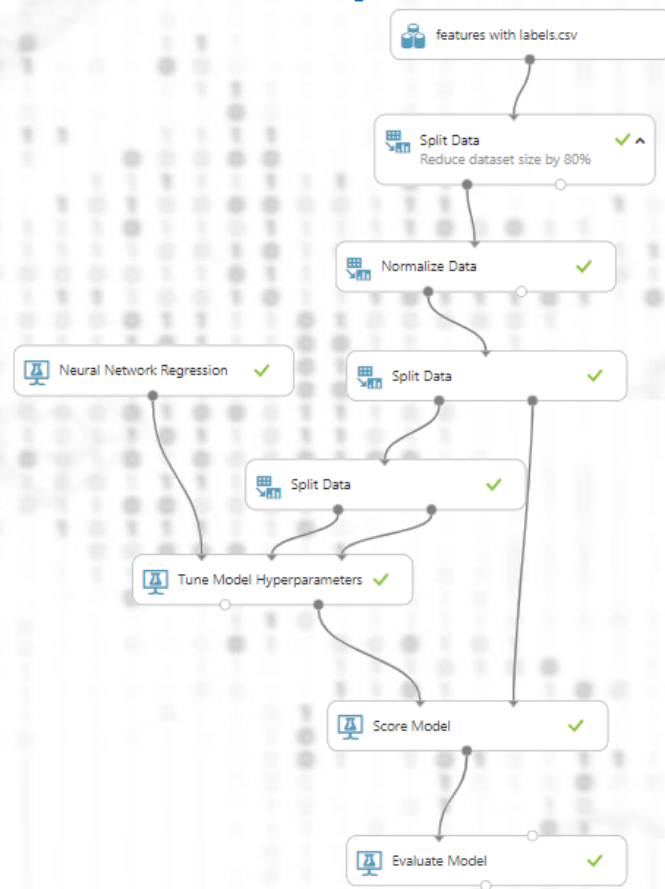
- ➔ Supervised (YOU provide the result)
 - ➔ Linear regression
 - ➔ Classification
- ➔ Unsupervised (YOU don't provide the result)
 - ➔ Clustering
 - ➔ Imaging (e.g. Neural networks)



I was told workflow was here. Could I see it, please?

Machine Learning typical workflow:

1. Get datasets
2. Clean, prep, feature engineering
3. Training (Choosing Algorithm and Hyper-Parameters)
 - Regression, Classification (2-class, multi-class), Clustering, Anomaly Detection, etc.
 - Deep Learning, Neural Networks, etc.
4. Scoring/Testing
5. Evaluating
6. Deploying predictive webservices (inference)
7. I'll retrain back!





Algorithm: Regression

➔ Learn by finding a linear pattern in data

➔ https://en.wikipedia.org/wiki/Linear_regression

➔ Predict numeric values
(linear or not)


➔ Trends

➔ Weather

➔ Economics


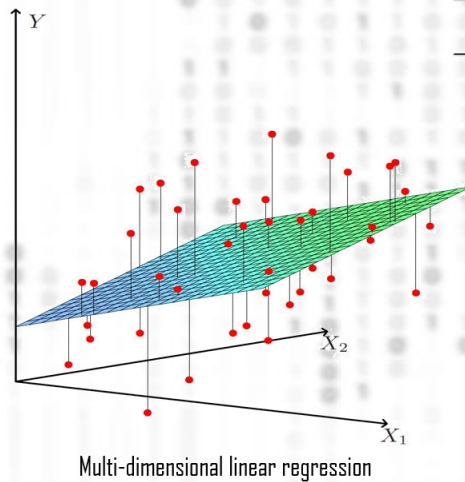

➔ Biology

➔ [...]








Algorithm: Regression






Algorithm: Regression


The inability for a machine learning method (like linear regression) to capture the true relationship is called **bias**.



Underfitted




Good Fit/Robust



Overfitted



Beware of the statistics monsters...






Algorithm: Logistic Regression

- Classify in 2-class model
- Learn by finding a binary pattern in data
 - Either the positive or the negative (or A and B)
 - Returns a numeric value (e.g. percentage, where 0% is most likely false and 100% is most likely true)

→ https://en.wikipedia.org/wiki/Logistic_regression





Data Scientist driven torture...

If you torture the data long enough,
it will confess.

Ronald Coase




online-behavior.com



Algorithm: Clustering




- ➔ Multiple algorithms for clustering
 - ➔ k-means, etc.
- ➔ Find grouping patterns in data
 - ➔ Shopping trends
 - ➔ Social network analysis
 - ➔ Crime analysis (fraud detection)



➔ https://en.wikipedia.org/wiki/K-means_clustering





Algorithm: Clustering with k-means





Algorithm: Clustering in N-dimensions




Explanation + code : <https://stanford.edu/~cpiech/cs221/handouts/kmeans.html>



Human statistics seen by aliens...

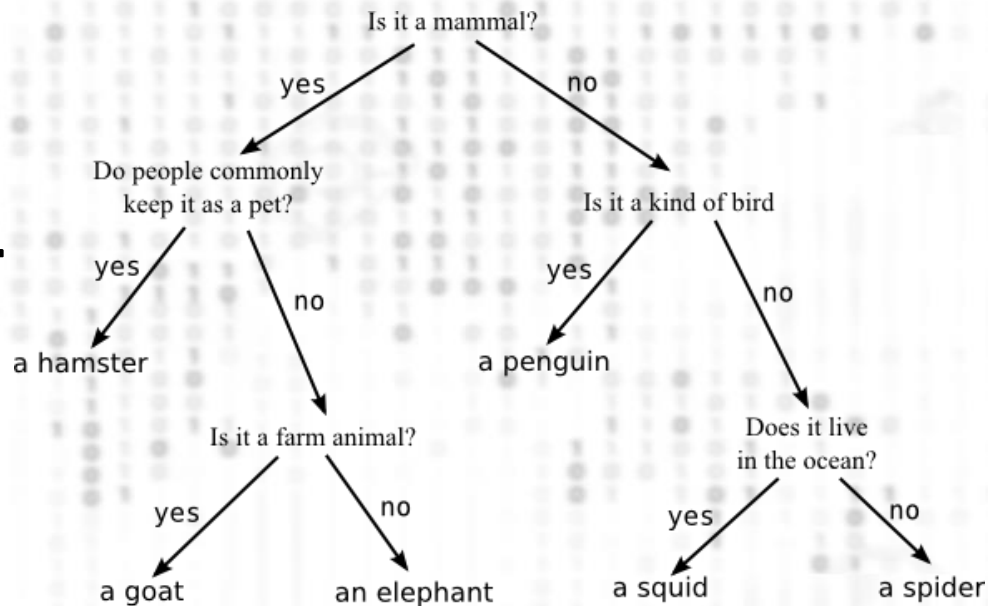
STATISTICS
THE DISCIPLINE THAT PROVES
THE AVERAGE HUMAN HAS
ONE TESTICLE





Algorithm techniques: Decision tree

- E.g. Multi-class model
- Both for numerical and categorical data.
- Conditions explained by simple Boolean logic.
- Easy to expand the tree if more detail is needed
- https://en.wikipedia.org/wiki/Decision_tree_learning





Decision tree in the wild (since the 1990s)

→ <http://20Q.NET>

Q3. Will it eat almost anything?

Yes , No , Unknown, Irrelevant, Sometimes, Maybe, Probably, Doubtful, Usually, Depends, Rarely, Partly

2. Can it scratch? **Yes.**

1. It is classified as **Animal**.



Play 20Q

About Us

Products

More ...

393

Like

Share

"The 20Q is so good at guessing, it's almost scary."
Stephen Cass
"IEEE Spectrum",

Q20. I am guessing that it is a kitty cat?
Right, Wrong, Close


19. Is it striped? **Sometimes.**
18. Is it brown? **Sometimes.**
17. Is it colourful? **Sometimes.**
16. Can it climb? **Yes.**
15. Can it growl? **Yes.**
14. Is it considered valuable? **Irrelevant.**
13. Would you give it as a gift? **Yes.**
12. Is it black? **Sometimes.**
11. Does it live in a burrow? **No.**
10. Does it dig holes? **Rarely.**
9. Is it killed for its fur? **No.**
8. Is it white? **Sometimes.**
7. Does it eat cheese? **No.**
6. Is it used with animals? **Irrelevant.**
5. Can you see it in a zoo? **Doubtful.**
4. Is it a wild animal? **Sometimes.**
3. Does it have teeth? **Yes.**
2. Is it small? **Yes.**
1. It is classified as **Animal.**




Algorithm techniques: Neural Networks

- E.g. Imaging, Deep Learning
- Slowest but most accurate
- https://en.wikipedia.org/wiki/Artificial_neural_network
- https://en.wikipedia.org/wiki/Deep_learning

Simple Neural Network



Deep Learning Neural Network



Input Layer



Hidden Layer

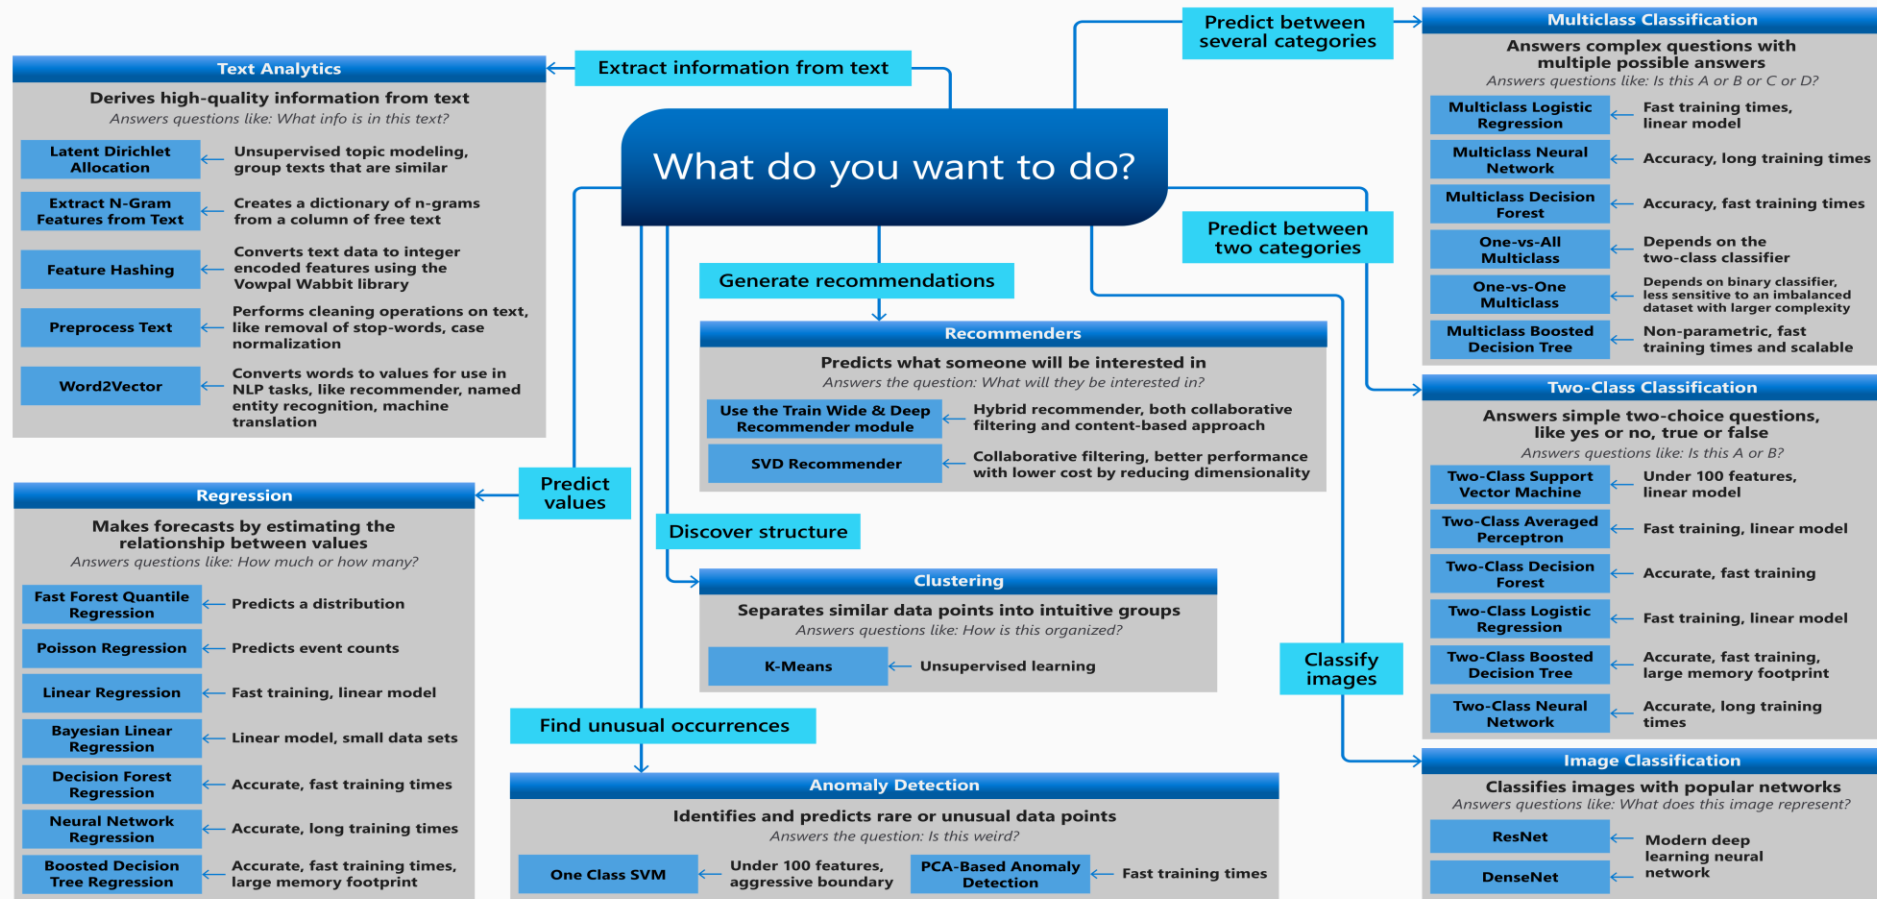


Output Layer



Machine Learning Algorithm Cheat Sheet

This cheat sheet helps you choose the best machine learning algorithm for your predictive analytics solution. Your decision is driven by both the nature of your data and the goal you want to achieve with your data.





Humans...

Using Machine Learning in the Microsoft
Universe doesn't require Data Science
background!


Anyone can do it!



"This is the world now. Logged on, plugged in, all the time."

1. Azure Machine Learning (Service / Workspace, not "Classic")
 - SDK Includes a lot of stuff (AutoML, HyperDrive, Model registry, SDK support for webservices in containers, etc.) – Supports Python, R (incomplete), MLFlow, etc.
 2. Azure Databricks (Spark clusters, SQL, Python, R, Scala)
 3. Azure HDInsight Spark
 4. Azure HDInsight R Server
 5. Azure Synapse Analytics (with SQL or Spark – SQL, Python, R, Scala, .Net)
 6. SQL Server 2016 (SQL, .Net or R Services)
 7. SQL Server 2017/2019/2022 (SQL, .Net, ML Services using R or Python)
 8. Any platform with .Net Core and libraries for ML.Net
 9. Azure Cognitive Services + Bots (managed webservices) + M365 Power Virtual Agents
 10. Etc...
- ❑ Some of these include their own notebooks (Jupyter, Zeppelin, DBC, etc.) and you can use your tools like Azure Data Studio, etc.

Magic is in the libraries,
not the languages!



"Come with me if you want to DEMO..."



Questions,
Humans?

ONE DOES NOT SIMPLY



UNPLUG SKYNET



"The unknown future rolls toward us.
I face it for the first time with a sense of
hope, because **if a machine, a Terminator,
can learn the value of human life,
maybe we can, too.**"

ধন্যবাদ

Vielen Danke!

Дуже дякую!

Благодаря!

Thank you!

Obrigado, pá!

¡Muchas gracias!

ধন্যবাদ্

Hvala vam!

මග්නාට ස්තූතියි

Ευχαριστώ!

Merci beaucoup!

Terima kasih!

Grazie mille!

Đakujem!

Mulțumesc!

Labai ačiū!

Dziękuję Wam!

Mockrát děkuju!

Mange tak!

Kiitos!

Takk fyrir!

Dank u wel!

Takk!

Dank je!

Tack så mycket!

Köszönöm!

Go raibh maith agaibh!

Diolch!



André Melancia



Andy.PT

LunarCat.PT