


DATA SCIENCE


15 0

07 00

THE
BIG DATA
CHALLENGE




THE
DATA
CHALLENGE



| | |
|------|------------|
| DATE | 15/07/2015 |
| TIME | |
| NAME | |

DANIELE
DANIELE QUERCIA
VICE



What is Big Data?

Huge, fast, and complex data

A Proper Definition. *"Data whose scale, diversity, and complexity require new architectures, techniques, algorithms, and analytics to manage it and extract value and hidden knowledge from it."*

Translation. There's too much data, coming in too fast, from too many places, and we need smarter ways to handle it.




Where is All This Data Coming From?

Web servers & log files. Everything you do online leaves a trace. Even that embarrassing search history.

Internet of Things (IoT). Smart devices, sensors, and even your fridge (yes, your fridge).

User-Generated Content. Social media posts, reviews, tweets, and every "fire emoji" comment under a selfie.



Why Should You Care?

Big Data is everywhere. It's behind Netflix recommendations, Google searches, and why your phone knows where you're going before you do. It's **reshaping business, healthcare, social media, and even emergency response.**

So, next time you wonder why YouTube suggests a conspiracy theory at 2 AM, just remember: Big Data is watching.

The VS of Big Data



Big Data

More Vs Than a Rap Battle

By now, you get it - Big Data is, well, BIG. But what makes it so tricky? Let's break it down with the 5 Vs that data scientists obsess over.

1. Volume. There's Just Too Much
Imagine trying to count every grain of sand on a beach... while more sand keeps getting dumped on your head. That's Big Data Volume—44x more data in 2020 than in '09.

2. Velocity. It Never Stops Coming
Data is generated constantly - live streams, stock prices, tweets, weather updates, and that one friend who texts "bro" 500 times a day.

The VS of Big Data



Big Data

More Vs Than a Rap Battle

3. Variety. It's a Mess. Numbers, Text, Audio, Video, Emojis 🤪. Big Data comes in all formats—structured, unstructured, or somewhere in between. Sorting it out is like trying to organize your school locker after ignoring it all year.

4. Veracity. *Can We Trust It?* Not all data is reliable. Some of it is gold, some is garbage, and some is as trustworthy as a clickbait headline. Data scientists have to clean and filter out the nonsense (like removing fake Google Maps reviews from people who clearly didn't even eat there).

The VS of Big Data


Big Data

More Vs Than a Rap Battle

5. Value. What's the Point? Data without purpose is like a Ferrari with no engine - useless. The goal is to turn all this data into something valuable - whether that's *improving healthcare, predicting earthquakes, or just making sure your Spotify playlist is decent.*



BIG DATA:



DATA SCIENCE
EXTRACTING MEANING
FROM LARGE DATA



What Is Data Science?


Let's be honest - data science sounds complicated, but it's really just about making sense of a ridiculous amount of data and using it to find hidden patterns.

The Official Definition. *"Extracting meaning from very large quantities of data."*

- D.J. Patil, the guy who actually coined the term data scientist.

How Do Data Scientists Do It?

Think of them as detectives - but instead of solving crimes, they solve problems using data. The process is called **Knowledge Discovery in Databases (KDD)** (yes, they love fancy names).



DATA SCIENCE
EXTRACTING MEANING
FROM LARGE DATA



The Data Science Process (KDD)


Here's how it works (1/2):

Generation. Data is collected from sensors, logs, and users (yes, even your Instagram posts).

Collection. The data is pulled from different sources (e.g., web crawling, clickstreams, surveillance). Sometimes people give it willingly (Google searches), sometimes they don't realize it's being taken (waves at cookies).

Transmission. It's sent to data centers at high speed (like a never-ending firehose of information).

Preprocessing. Cleaning the data, removing garbage, and making it useful (because raw data is a mess).



DATA SCIENCE
EXTRACTING MEANING
FROM LARGE DATA



The Data Science Process (KDD)

Here's how it works (2/2):

Storage. Huge databases hold everything (think of it as a supercharged version of your phone's storage, but infinite). Data is kept in massive databases using fancy tech like Hadoop, Cassandra, and MongoDB.

Analysis. This is where the magic happens. Statistical analysis, machine learning, and AI step in to find patterns, make predictions, and help businesses, scientists, and governments.




Data Mining. Real-World Examples

Supermarkets & Shopping Habits. Ever wonder why stores put chips near soda? Because data mining says you'll buy both.

Online Recommendations. Netflix, Amazon, and Spotify don't guess what you'll like—they use data mining to predict it.

Fraud Detection. Banks analyze patterns to catch criminals. (No, buying 10 PlayStations at 3 AM is not normal.)

PROFILING



Profiling





Yes, You're Being Analyzed

You leave a digital footprint everywhere: what you watch, what you buy, what you Google at 3 AM. Data profiling takes all that and turns it into insights that businesses, social media, and governments use (sometimes for good, sometimes... less so).

It's why Instagram knows you love sneakers before you even say it out loud.

PROFILING

Yes, You're
Being
Analyzed



Profiling

Where Does This Data Come From?


E-commerce Sites. Every product you click, add to cart, or abandon (we see you).

Search Engines. Your queries, searched topics, and even typos.

Social Networks. Profiles, posts, tweets, likes—every digital breadcrumb you leave behind.

Geolocation Data. Your phone tracks where you go (but sure, turn off location services, that'll fool them).

PROFILING



Yes, You're
Being
Analyzed

Profiling


Why Bother Profiling?

Companies want to sell you stuff, and governments want to understand trends.

Recommendation Systems. Netflix, Amazon, and Spotify tailor suggestions based on what you've engaged with.

Market Basket Analysis. Stores know if you buy diapers, you're probably also buying beer (new parents need a break).

Context-Aware Data Analysis. Ads changing based on your location, the time of day, and your browsing habits (why yes, I was just thinking about pizza).




Preprocessing Cleaning Up the Data Mess

Real-world data is dirty- messy, inconsistent, and full of errors. If you don't clean it up first, your analysis will be as useless as a broken calculator. That's where preprocessing comes in.

Why Bother?

Because bad data = bad results. Imagine a science experiment where half the measurements are wrong. Useless, right? Same with data science.



Preprocessing Why Does This Matter?



Without cleaning, your results are garbage.

Good data = good predictions. Whether it's healthcare, business, or science, clean data leads to better decisions.

Bottom Line? Data preprocessing is like washing vegetables before cooking—skip this step, and you're just serving up dirt.

THE REALITY OF DATA SCIENCE THE HARSH TRUTH

80% THE JOB IS SCIENCE



The Reality of Data Science

What Data Scientists Actually Do (80-90% of the Time):

Talk to Experts. You need to understand the domain before you can analyze the data.


Figure Out the Data Sources. Where's the data coming from? What does it mean?

Wrangle, Extract & Integrate. You have to dig through files, connect sources, and make it usable.

Clean the Data. Fix errors, remove duplicates, and sort out inconsistencies.

THE REALITY OF DATA SCIENCE THE HARSH TRUTH


80% OF THE JOB IS CLEANING AND PREPARING DATA, NOT MACHINE LEARNING



89%
IS THE CLEANING
AND PREPARING,
NOT MACHINE LEARNING



80%
OF THE CLEANING
AND PREPARING DATA
NOT MACHINE LEARNING



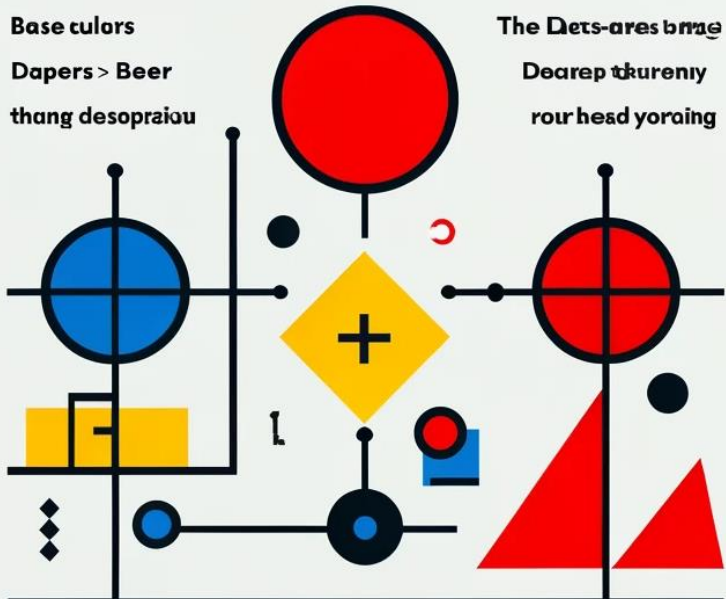
**80% PREPARING
DATA**

The Reality of Data Science

The Harsh Truth?

Machine learning is the glamorous part.
But 90% of the job is just making sure the data isn't a complete disaster.

ASSOCIATION RULES



WHEN DATA KNOWS YOU YOU THAN YOU



Eversfiengiitvele
unepesctiecting
data relationships



Association Rules

When Data Knows You Better Than You Do


Ever wonder why Amazon mysteriously suggests exactly what you need? That's **association rule mining** - finding relationships between items in data.

The Concept. It's all about spotting **patterns**. For example:

Diapers → **Beer** 😬

Turns out, a lot of people buying diapers also buy beer. Why? Stressed-out parents. Stores figured this out and started placing diapers near beer. Sales went up.

ASSOCIATION RULES



WHEN DATA KNOWS YOU YOU THAN YOU



Eversfiengiitvele
unepescfiecting
data relatorships




Association Rules

How It Works

- 1 Scan through millions of transactions.
- 2 Find patterns that happen often together.
- 3 Use those patterns to predict future behavior.

ASSOCIATION RULES



WHEN DATA KNOWS YOU BETTER THAN YOU



Eversfiengiitvele
unepescfiecting
data relationships



Association Rules


Real-World Examples

Netflix & Spotify. Recommending shows or songs based on what other people with similar tastes like.

Grocery Stores. Placing certain items together because you'll probably buy both.
Online Ads – That "random" ad for sneakers? Yeah, it's not random.

And yes, that's why Target somehow knew an adolescent was pregnant before she told anyone (even her parents).

Clustering: Finding groups



Clustering

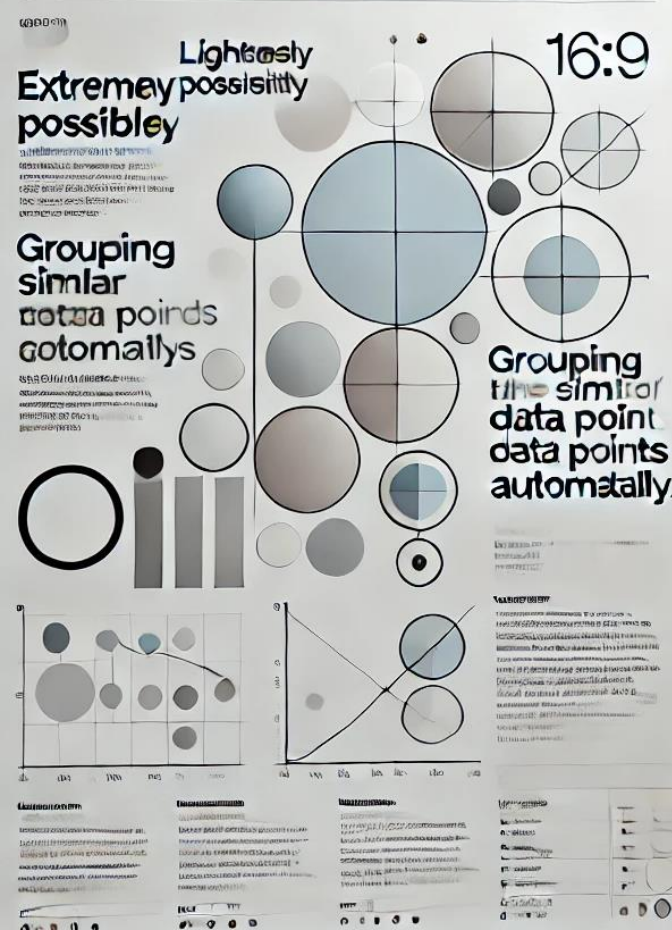
Finding Groups Without Even Trying

Clustering is like a school cafeteria—even without assigned seating, people naturally group together. Data does the same thing.

What's Clustering?

It's a technique that groups similar data points together - without knowing the groups in advance.

Clustering: Finding groups



Clustering


Why Does This Matter?

Retail & Marketing. Grouping customers based on their shopping habits (because not everyone wants glittery phone cases).

Medical Research. Identifying patterns in diseases or patient symptoms.

Streaming Services. Finding hidden user preferences (why else does Netflix magically know you love crime documentaries?)

Clustering: Finding groups



Clustering

Example: Spotting Outliers

Let's say 99% of people buy coffee and 1% buy 50 gallons of milk at once. Clustering can spot weird behaviors, like bulk buyers, potential fraud, or just someone with a really unhealthy dairy obsession.

Clustering helps us group things logically, even when we have zero idea what the groups should be. It's why Spotify knows whether you're more "chill beats".


CLASSIFICATION

Teaching Computers
to label stuff.

16:9

Classification:

Teaching computers how
to recognize for assigning labels
to data based patterns,
based patterns.



DERRECTY
DETTI
DUITE

DERRECTY
DETTI
YELLOW

Classification

Teaching Computers to Label Stuff

Classification is basically teaching computers how to sort things - like a super-smart digital sorting hat, but instead of putting you in Gryffindor, it decides if an email is spam.

What's Classification?


It's a technique that assigns labels to data based on patterns. The computer looks at past examples and learns to make future predictions.

CLASSIFICATION

Teaching a computer
to label stuff.

16:9

Classification:
Teaching profiles for
so teaching for asserting label
to data based patterns,
based patterns.



DERRECT
DEITTI
DUITE

DERRECTY
DEITTI
YELLOW

Classification

How It Works

Training Set. Feed the algorithm labeled examples (e.g., “this is spam”, “this is not spam”).

Learning Phase. The model figures out what features matter.

Test Set. We check if it got smarter or if it still thinks everything is spam.

Prediction. The final model labels new data on its own.

CLASSIFICATION

Teaching computers
to label stuff.

16:9

Classification:

Teaching computers to
so teaching for asserting label
to data based patterns,
based patterns.



Classification

Real-World Examples

Email Filters. Sorting spam from real emails (like saving you from that “Nigerian Prince” scam).


Facial Recognition. Identifying people in photos (yes, that’s why Facebook tags your friends before you do).

Medical Diagnosis. Classifying diseases based on symptoms (because “WebMD panic” isn’t a medical degree).

Artificial Neural Networks

TEACHING COMPUTERS TO THINK

teaching computers, sort of



NETURAL NETWORKS

influence recommendations

Teaching Neural Networks

BOATHOU BUBIF

AI

Artificial Neural Networks

Teaching Computers to Think (Sort Of)

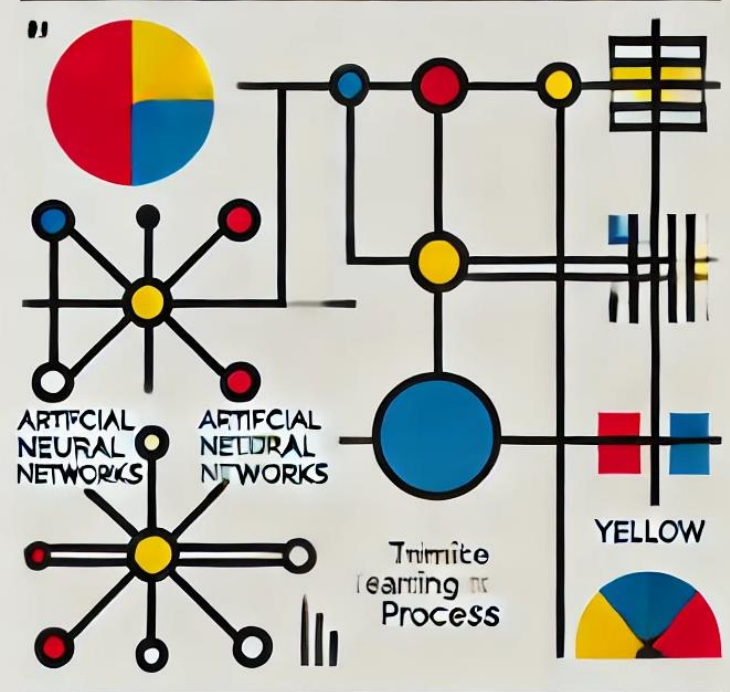
Ever wonder how AI recognizes faces, understands speech, or suggests the perfect meme? It's all thanks to Artificial Neural Networks (ANNs)—a system inspired by how your brain works, but without the bad decisions and procrastination.

And yes, this is why YouTube thinks you may really love conspiracy theories after watching just one video.

Artificial Neural Networks

TEACHING COMPUTERS TO THINK

teaching computers, sort of



Artificial Neural Networks

How It Works

Think of neurons in your brain: each one gets input, processes it, and sends a signal. Artificial Neural Networks do the same thing but with math instead of, well... brain cells.

Input Layer. Takes in raw data (like an image, text, or sound).

Hidden Layers. Where the magic happens! Neurons process the info and find patterns.


Output Layer. Gives a result (e.g., "Yes, that's a cat" or "No, that's just a weird-looking dog").

Artificial Neural Networks

TEACHING COMPUTERS TO THINK

teaching computers, sort of

| | | | |
|------------|------------|------------|------------|
| TECHNOLOGY | TECHNOLOGY | TECHNOLOGY | TECHNOLOGY |
|------------|------------|------------|------------|



NETURAL NETWORKS

influence recommendations

Teaching Neural Networks

BOATHOU BUTIP

AI

Artificial Neural Networks

Types of Neural Networks & What They Do


Feedforward Neural Networks (FFNNs). The basic model. Good at recognizing patterns. *Used in:* Predicting stock prices, recognizing handwritten text.

Convolutional Neural Networks (CNNs). Specializes in image recognition, so your phone knows it's you and not your boyfriend/girlfriend. *Used in:* Facial recognition, medical scans, self-driving cars.

Artificial Neural Networks

TEACHING COMPUTERS TO THINK

teaching computers, sort of



NETURAL NETWORKS

influence recommendations

Teaching Neural Networks

BOATHOU BUIP

AI

Artificial Neural Networks

Types of Neural Networks & What They Do

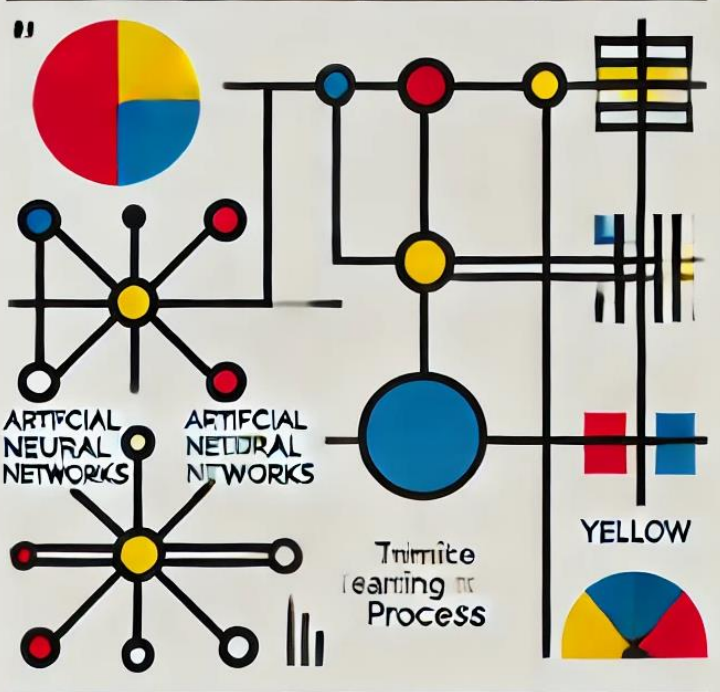
Recurrent Neural Networks (RNNs). Built for sequential data, meaning it actually remembers past inputs. *Used in:* Predicting text, speech-to-text, music generation.

Autoencoders. Great at removing noise from images, videos, and audio. *Used in:* Image enhancement, fixing blurry photos, compressing data.

Artificial Neural Networks

TEACHING COMPUTERS TO THINK

educating computers, sort of



Artificial Neural Networks

Real-World Uses


Facial Recognition. Your phone unlocks only for you (hopefully).

Self-Driving Cars. Identifies stop signs, pedestrians, and bad drivers.

Voice Assistants. Alexa, Siri, and Google listening in (don't pretend you didn't know).

Other AI Techniques:

Because Neural Networks Don't Run the whole Show



Sequence Mining
 Sequence Mining
 Sequence Mining
 Time Series Analysis
 Sequence Matrices Analysis
 Time Series Analysis
 Time Series Analysis
 Time Series Analysis
 Time Series Analysis
 Regression
 Time Series Analysis
 Outlier Detection
 Sequence Mining
 the show the show
 Outlier Regression
 Outlier Detection
 Association
 Regression

Other AI Techniques


Beyond Neural Networks

Sequence Mining. Finds patterns in ordered data. *Used in:* DNA analysis, shopping habits, and figuring out which Netflix show you'll binge next.

Time Series & Geospatial Data. Deals with data that changes over time or across locations. *Used in:* Weather forecasting, Google Maps traffic predictions, and making sure your Uber actually arrives.

Other AI Techniques:

Because Neural Networks Don't Run the whole Show



Other AI Techniques

Beyond Neural Networks

Regression. Predicts continuous values (instead of categories). **Used in:** Stock market predictions, house pricing, and guessing how long you'll scroll TikTok before realizing you should've been asleep hours ago.

Outlier Detection. Spots data that doesn't fit the pattern (a.k.a. weird stuff). *Used in:* Fraud detection, network security, removing outliers.

Sequence Mining the show

Outlier Regression

Outlier Detection

Association

Regression

THE DATA SCIENCE KITCHEN



The ~~data~~ Expert
Collects  data
Data Expert
Collects data









Data Anperst
Fonds Analyst
makes predecions



Data Analyst



Data Analyst   THE DATA 
  & MAKES PREDCTIONS


Finds patterns and
data ands visuals


Translaes
data in visuals

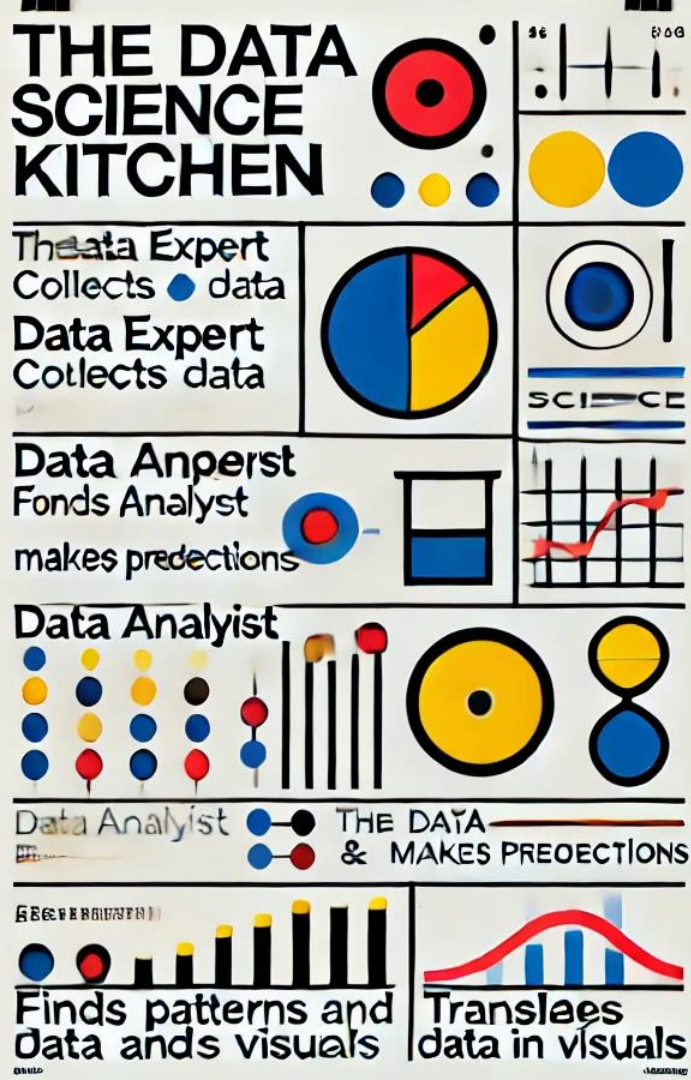
Data Science Roles

Data science takes a mix of skills and experts to actually make sense of all that data.

The Data Expert. Knows how to collect, process, and store massive amounts of data. Basically, the ingredient gatherer.

The Data Analyst. Uses statistics, machine learning, and data mining to find patterns and make predictions.

The Visualization Expert. Makes complex data look pretty and understandable (e.g., charts, graphs, dashboards).



Data Science Roles

The Domain Expert. Knows the actual field the data is from (e.g., healthcare, finance, social media) (Because data without **context** is just... numbers.)


The Business Expert. Uses data to make real-world decisions and create new business strategies. (aka the one who decides how much your next Uber ride is going to cost.)

Why Does This Matter?

Because data science isn't just "throw some numbers into a machine and hope for the best"- it takes teamwork, different skills, and actual human brains to make AI useful and trustworthy.

Open Issues:

Why Data Science Still isn't Perfect



Bias in Algorithms & Data

Interspetability & Transparency

Bias in Algorithms & Data

Orgsacy
Cnoens

Privacy
Privacy Conens

Privacy
Black Box
Problem

Black Box
Problem

Open Issues

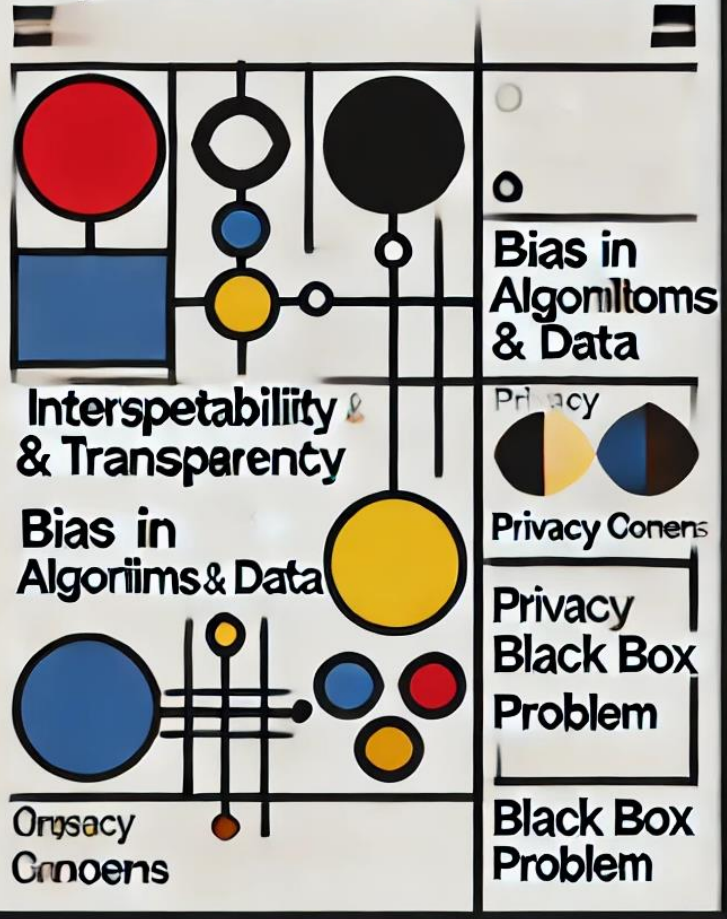
Why Data Science Still Isn't Perfect

Interpretability & Transparency. AI models make decisions, but sometimes even experts don't know how. It's like when your teacher marks something wrong, but their explanation is just "because I said so".

Bias in Algorithms & Data. AI is only as good as the data it learns from. If the data is biased, the AI will be too. *Example:* AI hiring tools that favored men because they were trained on old hiring data full of bias.

Open Issues:

Why Data Science Still isn't Perfect



Open Issues


Why Data Science Still Isn't Perfect

Privacy Concerns. Your data is valuable. Really valuable. Companies want it, hackers want it, and let's be real—you probably already gave it away by clicking 'Accept Cookies' without reading.

The Black Box Problem. Some AI models work, but no one really knows why. That's like trusting a self-driving car that won't tell you how it decides when to stop. (Slightly terrifying.)

Open Issues:

Why Data Science Still isn't Perfect



Bias in Algorithms & Data

Interspetability & Transparency

Bias in Algorithms & Data

Privacy
Privacy Consens

Privacy Black Box Problem

Orgsacy
Cnoens

Black Box Problem


Open Issues

Why Should You Care?

Because AI is influencing real-life decisions: hiring, healthcare, law enforcement. If we don't fix these issues, we could end up with AI making unfair or even dangerous choices.

INTERPRETABILITY IN MACHINE LEARNING

Model explanation
Prediction Selection



MODEL EXPLANATION
Model Explanation & Interpretable
Prediction Explanation
Prediction Feature selection

Interpretability in Machine Learning

AI is great at making decisions, but here's the problem - it often can't explain why it made them. Imagine your teacher giving you a bad grade and, when you ask why, they just say "can't tell you" That's AI for you.

Why Interpretability Matters

Model Explanation. We need to understand how an AI system works so we can trust it.

Prediction Explanation. If AI makes a decision, we should know why (especially if it affects jobs, loans, or medical diagnoses).

Interpretable Feature Selection. Choosing the right data matters.

