BENIN WORKSHOP ON ARTIFICIAL INTELLIGENCE

John AOGA

BIOGRAPHY

BWA'24 - Adversarial Search
25th May 2024

https://johnaoga.github.io/

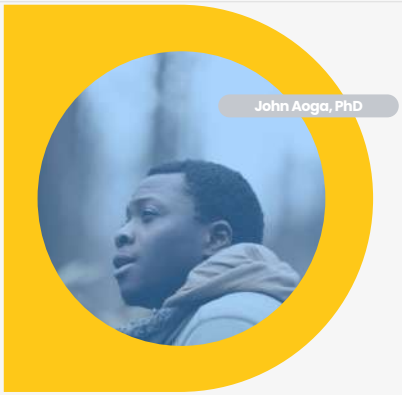

John Aoga, PhD

Who am I ?

Doctor & Engineer in Science and Technology
Specialist in Data science & AI
Researcher and Teacher
Co Founder of MIFY SARL company

Goals and Aspirations

Promote and develop AI 4 Africa In Africa
Promote and develop Education tools

Domains & Interests

Algorithms and Optimization
Data/Pattern Mining Approches and applications
Deep Learning & NLP for local languages
Social Data Analysis

Scientific References

Aoga John (johnaoga@gmail.com)

BENIN WORKSHOP ON ARTIFICIAL INTELLIGENCE

2

BENIN WORKSHOP ON ARTIFICIAL INTELLIGENCE

Adversarial search in AI

AI Beyond Learning

John Aoga

FRIARE

UCLouvain

Introduction

AI?

PROGRAMM
MACHINE
OBJECT

Reasoning
Simulate reasoning (understanding)

Adaptation
Can understand & live in its environment

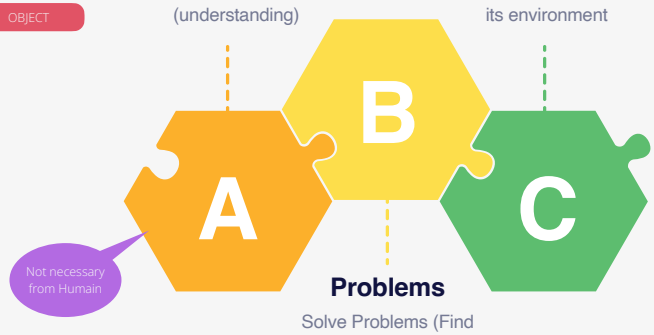

Problems
Solve Problems (Find feasible solutions)

INTELLIGENT AGENT

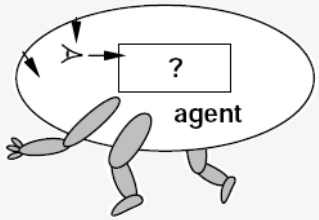

INTELLIGENT AGENT

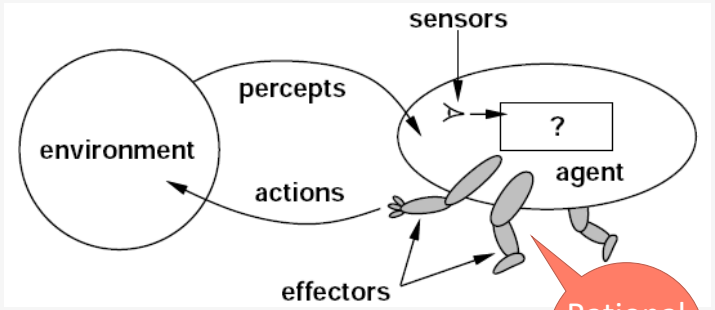

Rational actions

INTELLIGENT AGENT RATIONALITY

Need of search

Smart Vacuum

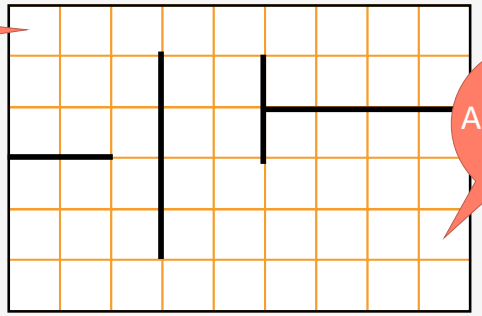

Adversity!

SEARCH IN AI

Type of search

(2) Informed Search

Uses problem-specific knowledge or heuristics to find solutions more efficiently - A*/ Hill Climbing - Route finding and Navigation

(3) Local Search

Iteratively explores the search space by moving to neighboring states - tabu search - traveling salmon problem

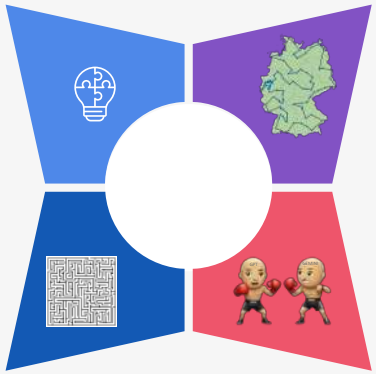

(1) Uninformed

Brute-force , explore all the search space blindly (without any knowledge about the problem) - DFS/BFS - Maze solving

(4) Adversarial Search

algorithmic approach in AI where agents compete against each other, aiming for optimal outcomes in games or simulations - Minimax/Alpha-Beta -

* There are more categories for example: constraint programming, metaheuristics search, Stochastic search, ...

Aoga John (johnaoga@gmail.com)

Stories

A.I. TIMELINE




1950 — **TURING TEST**
Computer scientist Alan Turing proposes a test for machine intelligence. If a machine can trick humans into thinking it is human, then it has intelligence

1955 — **A.I. BORN**
Term 'artificial intelligence' is coined by computer scientist, John McCarthy to describe "the science and engineering of making intelligent machines"

1961 — **UNIMATE**
First industrial robot, Unimate, goes to work at GM replacing humans on the assembly line

1964 — **ELIZA**
Pioneering chatbot developed by Joseph Weizenbaum at MIT holds conversations with humans

1966 — **SHAKEY**
The 'first electronic person' from Stanford, Shakey is a general-purpose mobile robot that reasons about its own actions

A.I. WINTER
Many false starts and dead-ends leave A.I. out in the cold

1997 — **DEEP BLUE**
Deep Blue, a chess-playing computer from IBM defeats world chess champion Garry Kasparov

1998 — **KISMET**
Cynthia Breazeal at MIT introduces KISmet, an emotionally intelligent robot insofar as it detects and responds to people's feelings








1999 — **AIBO**
Sony launches first consumer robot pet dog AiBO (AI robot) with skills and personality that develop over time

2002 — **ROOMBA**
First mass produced autonomous robotic vacuum cleaner from iRobot learns to navigate and clean homes

2011 — **SIRI**
Apple integrates Siri, an intelligent virtual assistant with a voice interface, into the iPhone 4S

2011 — **WATSON**
IBM's question answering computer Watson wins first place on popular $1M prize television quiz show *Jeopardy*

2014 — **EUGENE**
Eugene Goostman, a chatbot passes the Turing Test with a third of judges believing Eugene is human

2014 — **ALEXA**
Amazon launches Alexa, an intelligent virtual assistant with a voice interface that completes shopping tasks

2016 — **TAY**
Microsoft's chatbot Tay goes rogue on social media making inflammatory and offensive racist comments

2017 — **ALPHAGO**
Google's A.I. AlphaGo beats world champion Ke Jie in the complex board game of Go, notable for its vast number (2^{170}) of possible positions

MORE RECENT TIMELINE

Where AI is heading?

GPT-1, GPT-2, **AlphaStar** (Grandmaster level at StarCraft II)

ChatGPT, Microsoft/Github Copilot

2017 — **AlphaGo**, **AlphaZero**, Transformer arch.

2019

2020 — GPT3, AlphaFold,

2022

2023 — Gemini, Bard, GPT-4,

Aoga John (johnaoga@gmail.com)

SEARCH IN AI

The search space of problem to solve is very big

Chess

10^120

10^170

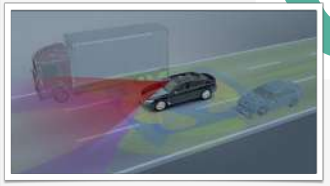

10^40 = 10000000000000000000000000000000000000000

Aoga John (johnaoga@gmail.com)

Adversarial Search Basics

TIC-TAC-TOE

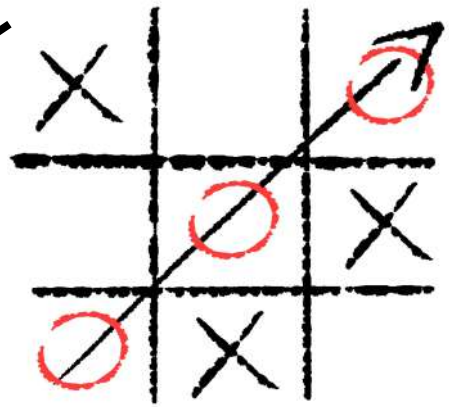

https://editor.p5js.org/yuvalbi/sketches/HuinyONTR

TIC-TAC-TOE SIMULATION

How to find the right action to take?

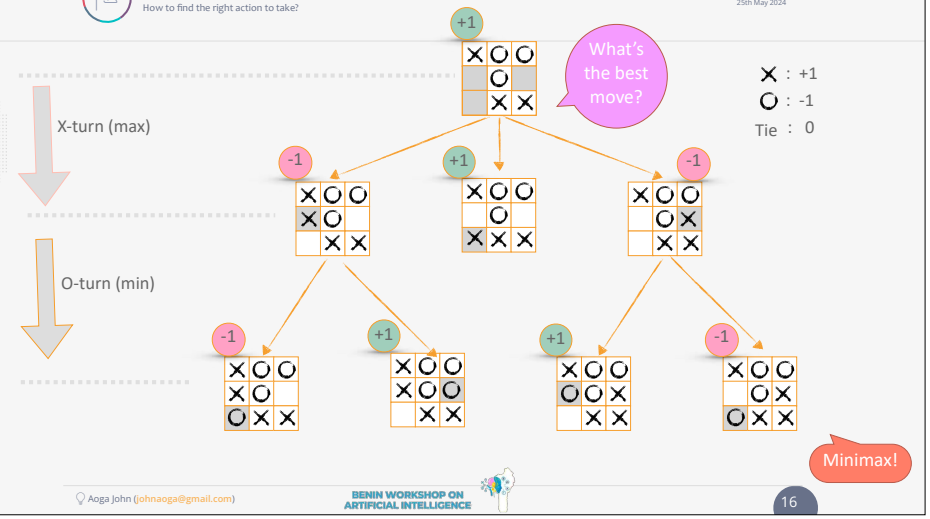

Aoga John (johnaoga@gmail.com)

Minimax

MINIMAX ALGORITHM SIMULATION

Simple example

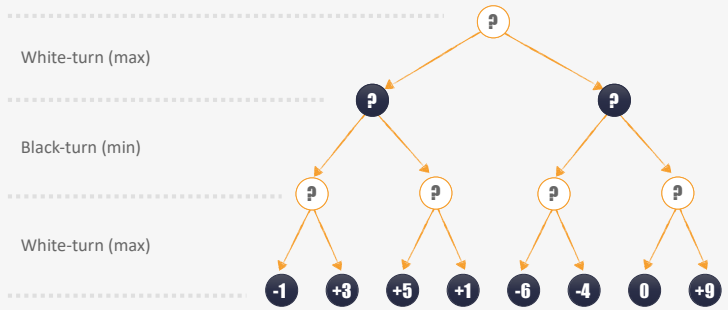

MINIMAX ALGORITHM IMPLEMENTATION

Simple example


```
function minimax(position, depth, maximizingPlayer)
    if depth == 0 or game over in position
        return static evaluation of position

    if maximizingPlayer
        maxEval = -infinity
        for each child of position
            eval = minimax(child, depth - 1, false)
            maxEval = max(maxEval, eval)
        return maxEval

    else
        minEval = +infinity
        for each child of position
            eval = minimax(child, depth - 1, true)
            minEval = min(minEval, eval)
        return minEval
```

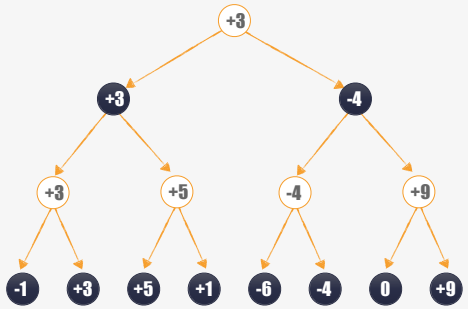

TIC-TAC-TOE SIMULATION

How to find the right action to take?

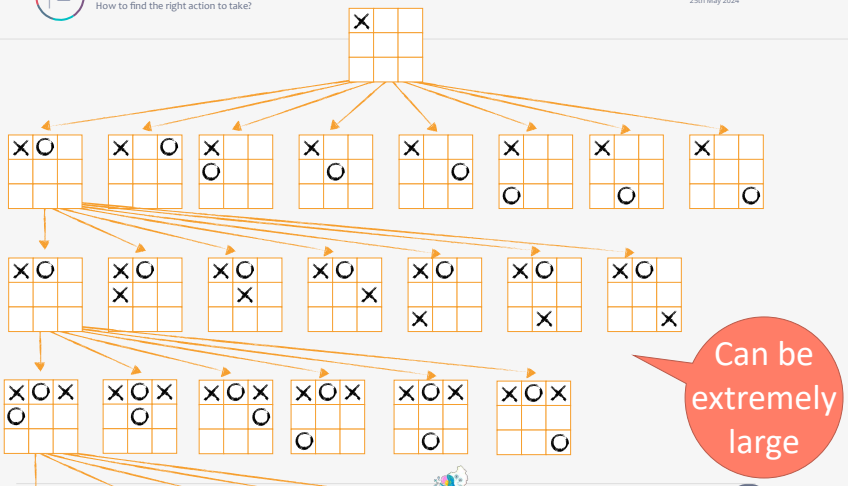

Alpha-Beta Pruning

ALPHA-BETA ALGORITHM

Simple example

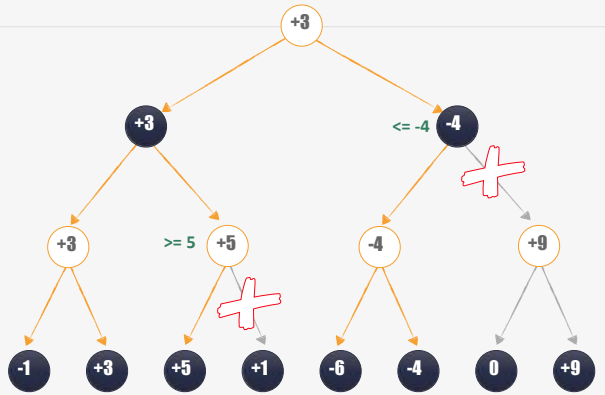

ALPHA-BETA ALGORITHM SIMULATION

Intermediate example

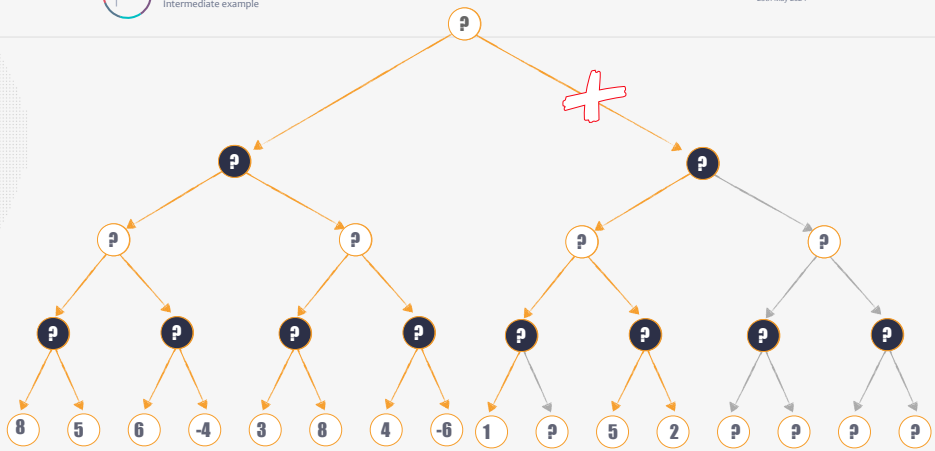

ALPHA-BETA ALGORITHM IMPLEMENTATION

Intermediate example

```
function minimax(position, depth, alpha, beta, maximizingPlayer)
  if depth == 0 or game over in position
    return static evaluation of position

  if maximizingPlayer
    maxEval = -infinity
    for each child of position
      eval = minimax(child, depth - 1, alpha, beta, false)
      maxEval = max(maxEval, eval)
      alpha = max(alpha, eval)
      if beta <= alpha
        break
    return maxEval

  else
    minEval = +infinity
    for each child of position
      eval = minimax(child, depth - 1, alpha, beta, true)
      minEval = min(minEval, eval)
      beta = min(beta, eval)
      if beta <= alpha
        break
    return minEval
```

Monte Carlo Tree Search

MCTS
Steps

BWA'24 - Adversarial Search
25th May 2024

Expansion

Expand the tree by adding new possibilities that stem from the current game state represented by the node

Simulation

From the new node, simulate a random playout to the end of the game.

SELECTION

Starting from the root node, select successive child nodes down to a leaf node - UCB1

Backpropagation

Propagated back up the tree, updating the statistics of the nodes visited during the selection phase.

Ethics Implications

AI as a threat to our Planet ?!

Three mains fields in NLP

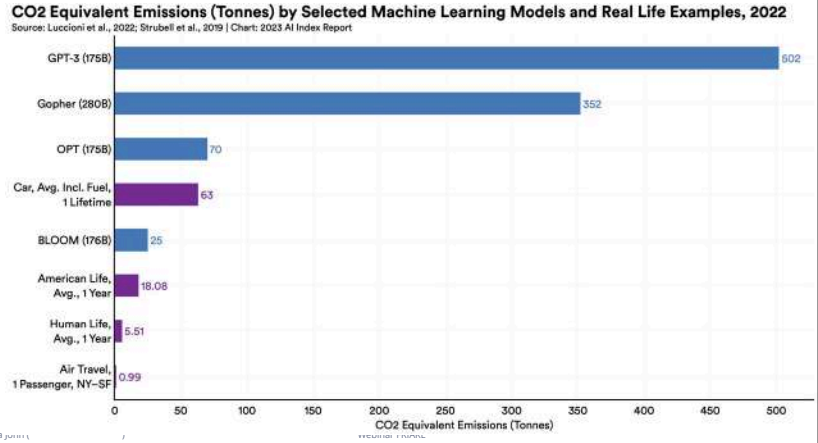

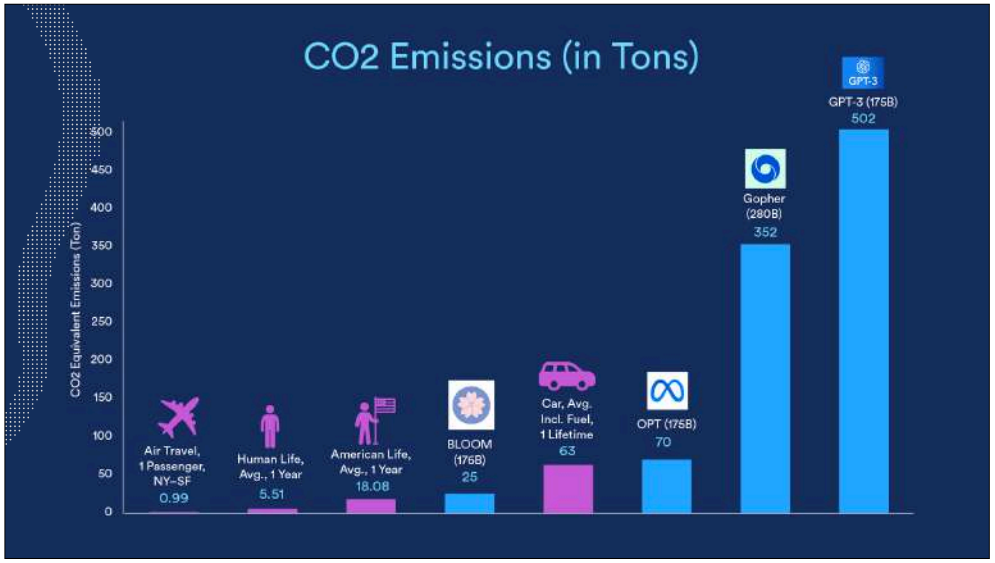
CO2 Emissions (in Tons)
CO2 Equivalent Emissions (Ton)
0
50
100
150
200
250
300
350
400
450
500
Air Travel, 1 Passenger, NY-SF
0.99
Human Life, Avg., 1 Year
5.51
American Life, Avg., 1 Year
18.08
BLOOM (176B)
25
Car, Avg. Incl. Fuel, 1 Lifetime
63
OPT (175B)
70
Gopher (280B)
352
GPT-3
GPT-3 (175B)
502

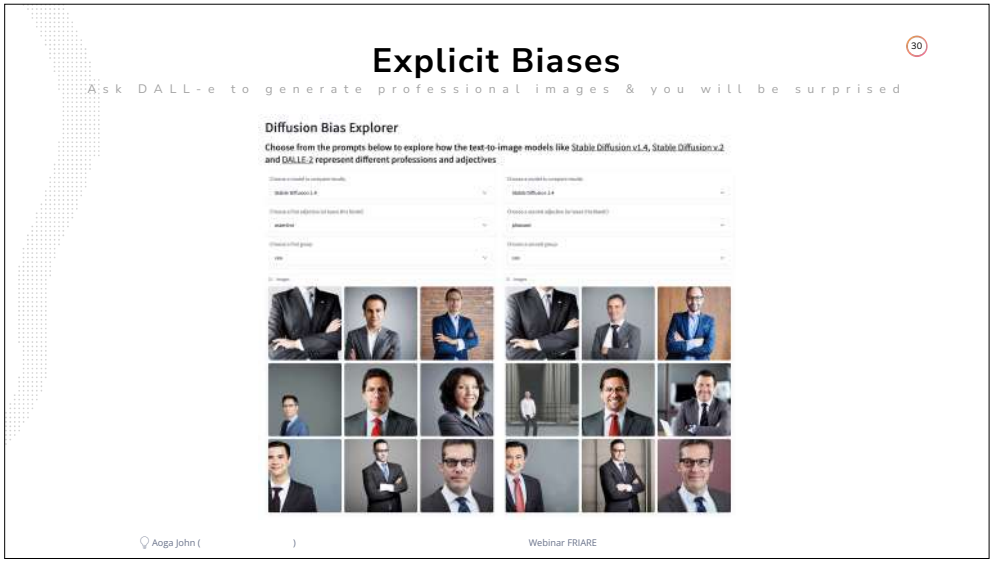
30
Explicit Biases
Ask DALL-e to generate professional images & you will be surprised
Diffusion Bias Explorer
Choose from the prompts below to explore how the text-to-image models like Stable Diffusion v1.4, Stable Diffusion v.2 and DALLE-2 represent different professions and adjectives
Aoga John (
)
Webinar FRIARE

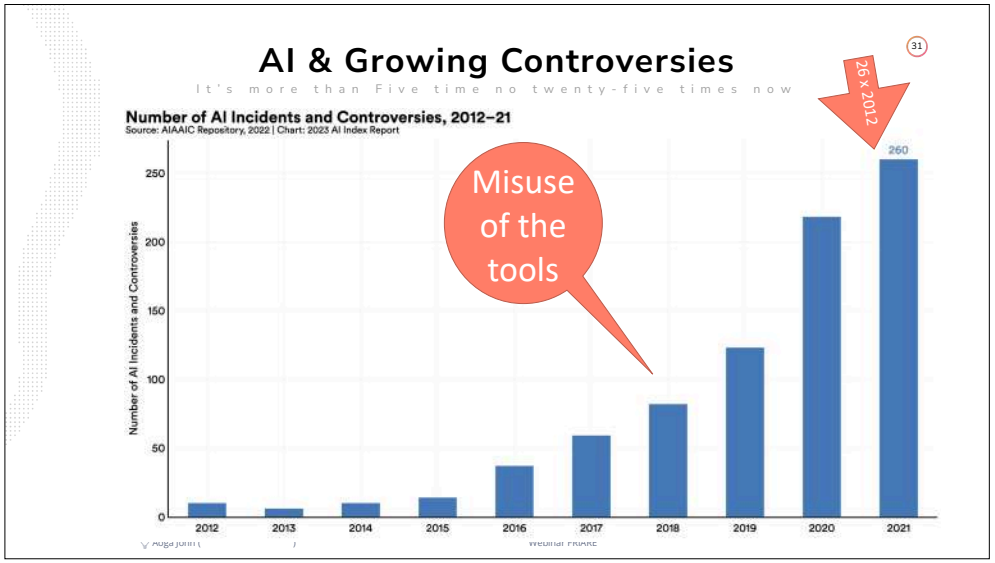
31
AI & Growing Controversies
It's more than Five time no twenty-five times now
Number of AI Incidents and Controversies, 2012–21
Source: AIAAIC Repository, 2022 | Chart: 2023 AI Index Report
26 x 2012
260
Misuse of the tools
Number of AI Incidents and Controversies
0
50
100
150
200
250
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021


Takeaway Message

Bible?!

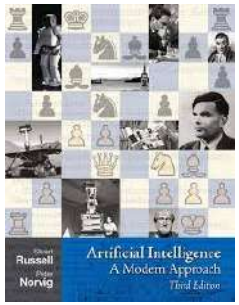

ANY QUESTIONS?