


Prof. Matthew England
Coventry University

The AI Landscape in 2026

SCOPING AI LITERACY AND READINESS OF WOMEN
ENTREPRENEURS IN KENYA AND INDIA

Stakeholders Policy Workshop
16th July 2026




This talk

This talk:

- Does not assume any prior knowledge of AI.
- Will not cover technical details.
- Will try to give an overview of:
 - What AI is in 2026
 - The landscape of different types of AI
 - Some key questions to ask when starting with AI.
 - What the latest developments are.

Do feel free to ask questions as we proceed!




Contents

1. What is AI?


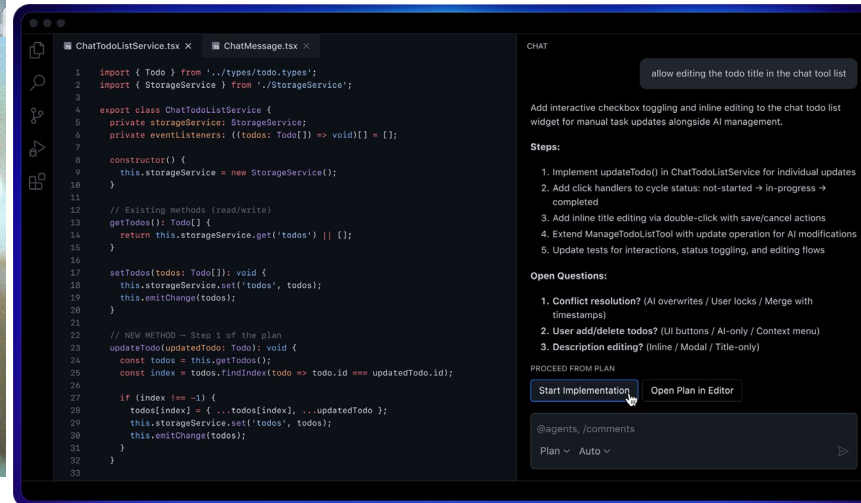


2. Types of AI/ML.

3. Key Questions for AI Beginners to Ask.

4. What's Changed Recently?




What is AI?



What is AI Technically?

Artificial Intelligence (AI): *computer systems designed to perform tasks that typically require human intelligence.*




What is AI Technically?

Artificial Intelligence (AI): *computer systems designed to perform tasks that typically require human intelligence.*

The definition is not very good because those tasks that “*typically require human intelligence*” is something that changes over time. E.g.

- Once, all arithmetic calculations were done by humans, then we invented calculators.




What is AI Technically?

Artificial Intelligence (AI): *computer systems designed to perform tasks that typically require human intelligence.*

The definition is not very good because those tasks that “*typically require human intelligence*” is something that changes over time. E.g.

- Once, all arithmetic calculations were done by humans, then we invented calculators.
- Once, all route planning was done by humans, but now we have Google Maps.



What is AI Technically?

Artificial Intelligence (AI): *computer systems designed to perform tasks that typically require human intelligence.*

The definition is not very good because those tasks that “*typically require human intelligence*” is something that changes over time. E.g.

- Once, all arithmetic calculations were done by humans, then we invented calculators.
- Once, all route planning was done by humans, but now we have Google Maps.
- Once, we assumed an essay could only be written by a human, then ChatGPT.

It is more useful to think about the underlying technology.




ChatGPT




How do we achieve AI?

We first sought to achieve AI by writing clever algorithms and data-structures for knowledge, with some success (AI beat humans at Chess in 1997 and Jeopardy in 2011)



```
while (user != "bye"):
    keyword_found = False


    for index in range(len(keywords)):
        if (keywords[index] in user):
            print("Bot: " + responses[index])
            keyword_found = True
```



Machine Learning

But AI in 2026 is dominated by one type of technology:

Machine Learning (ML): statistical algorithms to build models that learn patterns in data, and generalise to unseen data *without explicit instructions*.






Machine Learning

But AI in 2026 is dominated by one type of technology:

Machine Learning (ML): statistical algorithms to build models that learn patterns in data, and generalise to unseen data *without explicit instructions*.

Key ingredients of ML:

- Statistical Models which are formed by training on data
- The most prominent such model is the **deep neural network**, (very loosely) modelled on the brain.
- Lots of computing power and **lots of data** to train.




Machine Learning underpins modern AI

Machine learning underpins ChatGPT and other AI tools that have exploded in the last few years (as we'll see later). However, it also underpins much of the modern technology of the past 20 years. E.g.

- **Recommendation Systems:** YouTube recommends you videos, Netflix shows, Amazon products, all using machine learning models trained on their user data.
- **Autonomous Driving Systems:** Your car recognises road signs with models trained on driving data
- **Fraud Detection:** Banks detect fraudulent payments via machine learning models trained on their records.



The Netflix logo, featuring the word "NETFLIX" in a bold, red, sans-serif font, centered within a white rounded rectangle.





Why now? Big Compute, Big Data (and big money)


Moore's Law: The number of transistors on microchips doubles every two years

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important for other aspects of technological progress in computing – such as processing speed or the price of computers.


Transistor count




History of Machine Learning Models



McCulloch-Pitts




Rosenblatt




Widrow-Hoff


Minsky-Papert




Rumelhart, Hinton et al.



LeCun



Hinton-Ruslan Krizhevsky et al.



Vaswani

<https://medium.com/@lmpo/a-brief-history-of-ai-with-deep-learning-26f7948bc87b>


Contents

1. What is AI?

2. Types of AI/ML.

3. Key Questions for AI Beginners to Ask.

4. What's Changed Recently?





Types of Machine Learning I

Supervised Learning: identify patterns based on existing large labelled datasets.

May be used to solve:

- **Classification** Problems: E.g.,
“Does this image of a lung have cancer?”
(Based on this dataset of labelled images of lungs).
- **Regression** Problems: E.g.
“How much do you think this house is worth?”
(Based on this spreadsheet of houses and their features).




Types of Machine Learning II

Unsupervised Learning: identify patterns in large unlabelled datasets (this is much harder).

May be used to solve:

- **Clustering** Problems. E.g.,
“Can we split our customers into groups for marketing?” (Based on records of purchasing habits).
- **Dimensionality** Reduction. E.g.,
“Instead of using these 100 features about movies can we identify the three most significant to allow users to search by?”




Types of Machine Learning III

Reinforcement Learning: in this case there is no dataset; a machine learning model is trained by interaction with an environment which gives it rewards or penalties.


May be used to teach AI:

- **Games:** E.g.,
“Learn how to win at Chess/Go/...
(by playing millions of games with yourself).”
- **To Control Robots:** E.g.
“Learn how to make this robot move from here to there (by trying multiple times and being rewarded by how close you get).”





**There is a
LOT more
machine
learning
out there!**



So where do LLMs fit in?



Large Language Models (LLMs) have at their core a supervised learning model to predict the next word in a sentence.

E.g. given the text “the cat sat on the” then a reasonable prediction for the next word is “mat”.

In this example, an LLM answers “mat” not because it has been programmed to know anything about cats (or even about English grammar). It answers mat because this is statistically likely within the training data.

The training data for LLMs is the public internet!

Later, LLMs also benefitted from reinforcement learning (originally by teams of humans) to teach them to be helpful and safe.




Generative AI


LLMs have inspired a new category of AI/ML: **Generative AI**: uses ML technology to create “*new*” content.

LLMs create text. But similar technology has been built to create images, music, video, code, voices and more!

Now often packaged together into a single multi-modal AI tool.




	Traditional Analytical ML	Generative AI
Core Objective	Predict & Classify new data based on patterns in training data.	Generative: To generate new objects based on patterns in training data
Input Data Type	Traditionally data is structured (e.g. spreadsheet)	Designed for unstructured data
Output Data Type	A label, score, class, or numeric value.	A new object, e.g. text, image, code.
Underlying technology	A wide variety to choose from which vary in power, computational requirements, data requirements and interpretability.	Deep neural networks (Transformer-based architectures).
Scope	Specialists: Model is usually trained for just one task.	Generalists: The same model can be applied to many different tasks.
Explainability	Low-High: Depending on the underlying technology.	Low: Although you can ask an LLM to explain its reasoning this is not guaranteed to be accurate!
Accuracy	High: Easy to evaluate, and if trained correctly, they yield consistent, reliable results	Variable: Hard to evaluate given generality; susceptible to "hallucinations" (confident but false assertions).
Computational Cost	Low to moderate: many models can be trained on a standard laptop.	Massive: Requires hundreds of GPUs, millions of dollars, and massive power consumption. Commonly obtained as a cloud service.



Frontier Models Quiz





ANTHROPIC






ANTHROPIC






**Myth: All AI is
made in the USA**








**Myth: All AI is
made in the USA**



Mistral AI




Myth: All AI is now Generative AI



2024 Nobel Prize for Chemistry went to DeepMind for AlphaFold.

Supervised learning with specialised deep neural networks.




Contents

1. What is AI?

2. Types of AI/ML.

3. Key Questions for AI Beginners to Ask.

4. What's Changed Recently?



Key questions for beginners in AI 1:

(Q) Do you really need AI? AI is great for spotting patterns that you cannot; but if you already know the patterns, then traditional *digitalisation* and *automation* may be what you need!



Key questions for beginners in AI 1:

(Q) Do you really need AI? AI is great for spotting patterns that you cannot; but if you already know the patterns, then traditional *digitalisation* and *automation* may be what you need!

Note: In the past you would have had to invest a lot for automation, but LLMs can now help you produce code at a fraction of the cost.



Key questions for beginners in AI 2:

(Q) Generative AI or Traditional Analytical ML?

If you have lots of labelled structured data and you want to extract value from it, then you can probably do this using a traditional *supervised machine learning tool* instead of an LLM. This has the advantages of being:

- lower cost;
- more reliable results; and
- more interpretable.




Key questions for beginners in AI 3:

(Q) Is your data of high enough quality?

AI runs on data.

If using a generative AI tool, then the data was the – does that make sense for your application?

If your own data, then is the data representative?



Important: bias in data → bias in AI

Dataset trained Microsoft chatbot to spew racist tweets

In March 2016, Microsoft learned that using Twitter interactions as training data for ML algorithms can have dismaying results.


Within 16 hours, the chatbot posted more than 95,000 tweets, and those tweets rapidly turned overtly racist, misogynist, and anti-Semitic. Microsoft quickly suspended the service for adjustments and ultimately pulled the plug.

Amazon AI-enabled recruitment tool only recommended men

Like many large companies, Amazon is hungry for tools that can help its HR function screen applications for the best candidates. In 2014, Amazon started working on AI-powered recruiting software to do just that. There was only one problem: The system vastly preferred male candidates. In 2018, Reuters broke the news that Amazon had scrapped the project.



AI algorithms identify everything but COVID-19

Derek Driggs, an ML researcher at the University of Cambridge, together with his colleagues, published a paper in *Nature Machine Intelligence* that explored the use of deep learning models for diagnosing the virus. The paper determined the technique wasn't fit for clinical use. For example, Driggs' group found that their own model was flawed because it was trained on a data set that included scans of patients that were lying down while scanned, and patients that were standing up. The patients who were lying down were much more likely to be seriously ill, so the algorithm learned to identify COVID risk based on the position of the person in the scan.




Key questions for beginners in AI 4:

(Q) How important is AI model interpretability to you? There are a wide variety of AI models out there, some can be interpreted by humans more easily than others.



There is a perceived trade-off between accuracy and interpretability of AI.




Key questions for beginners in AI 4:

If using an LLM, then do you need a proprietary model through the cloud, or could a small open-weight model work?

Enterprise Cloud AI incurs a variable cost (based on size of input and output), and may rise in the future.

There do exist open-weight models which perform almost as well!




Key questions for beginners in AI 5:

How to ensure your AI project meets data security requirements.

If you use a free cloud AI tool then your data is certainly being used for future training.

Personal pro accounts do not usually meet UK / EU legal requirements for personal data security (need enterprise licences).

The EU has laws such as “the right to be forgotten” meaning people can ask for their data to be deleted. If that data trained an ML model then the data is encoded in that – how to remove?



What is an AI Hallucination?

Basically, a hallucination is when the text generated by an LLM is false but confident.

Sidenote: I really dislike the term “Hallucination”!


The text generated is always statistically likely based on the training data. It is astonishing that statistical likelihood takes us this far! But it is not, and never can be, a guarantee of correctness.

If you want AI that is **guaranteed** correct, then it cannot be solely machine learning based.



Compared to a couple of years ago:

- LLMs making far fewer hallucinations.
- It is now easier to identify hallucinations: request the sources of facts and then **check them!**
Note that the LLM itself is rarely critical of its sources: you need to do that critical review!
- New tools being developed that focus LLMs on “high quality” or “user selected” data.
New techniques being developed to minimise hallucinations.
- But there are still hallucinations there, and they cannot be removed completely.



Key questions that ~~beginners~~ everyone should ask:

(Q) How will I judge whether my AI produced content is correct?

(Q) What risk mitigations do I have in place for the cases where it is not correct?




Contents

1. What is AI?

2. Types of AI/ML.

3. Key Questions for AI Beginners to Ask.

4. What's Changed Recently?




From Chatbots to Agents

LLMs can now plan multi-step workflows and execute them using tools:

- They can search the web
- They can use software
- They can write and execute code
- They can read and write to files
- They can coordinate with other agents

This is increasingly being considered a new AI category of **Agentic AI**.




How to improve my LLM Output?


- Use the latest LLMs.
- Experiment with different LLMs.
- Experiment with prompt engineering techniques.



Prompt Engineering Techniques




<https://medium.com/the-modern-scientist/prompt-engineering-classification-of-techniques-and-prompt-tuning-6d4247b9b64c>



How to improve LLM Output?


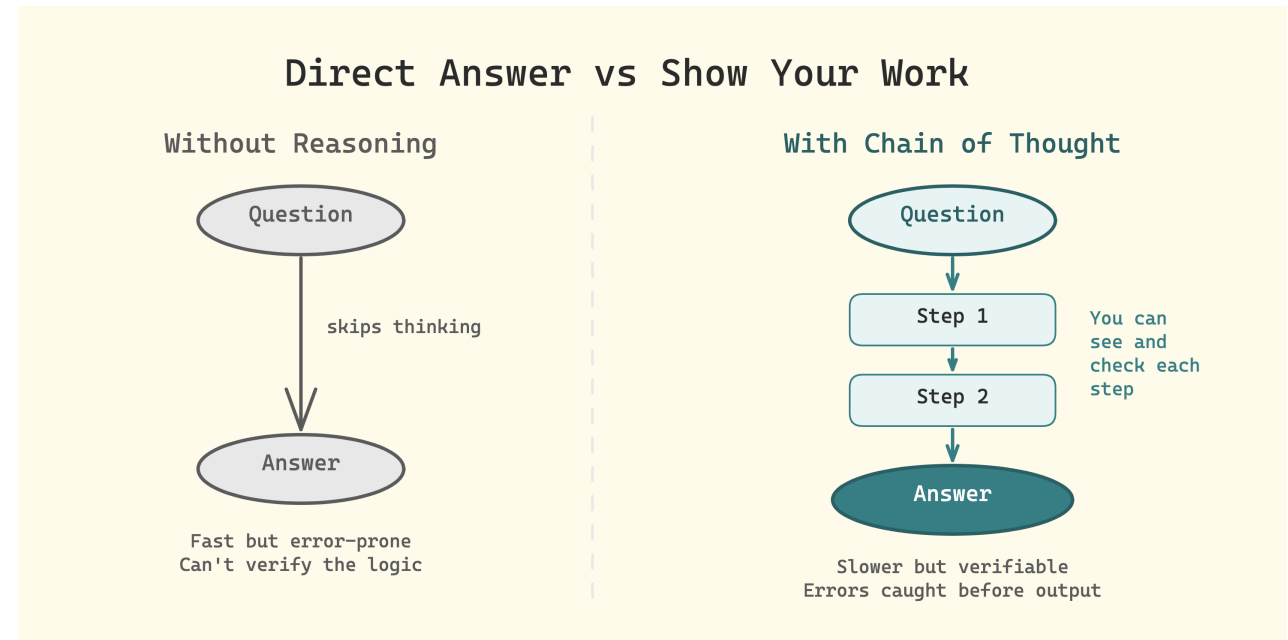
- Use the latest LLMs.
- Experiment with different LLMs.
- Experiment with prompt engineering techniques.
- Use a RAG system.

Retrieval Augmented Generation



How to improve LLM Output?

- Use the latest LLMs.
- Experiment with different LLMs.
- Experiment with prompt engineering techniques.
- Use a RAG system.
- Use a CoT or MoE model.



How to improve LLM Output?

- Use the latest LLMs.
- Experiment with different LLMs.
- Experiment with prompt engineering techniques.
- Use a RAG system.
- Use a CoT or MoE model.
- Combine LLMs with other tools.


Plot two sine curves with a frequency ratio corresponding to a perfect third. Show them in the colors of the Portuguese flag.

A perfect third in music consists of two frequencies with a ratio of 5:4. To plot two sine curves with this frequency ratio and the colors of the Portuguese flag (green and red), you can use the following code:

```
Plot[
  {Sin[5 x], Sin[4 x]},
  {x, 0, 4 Pi},
  PlotStyle -> {Darker[Green], Red},
  AxesLabel -> {"x", "Amplitude"},
  PlotLegends -> {"Frequency: 5x", "Frequency: 4x"}
]
```


> Used Wolfram Language Evaluator 

Here's the plot of two sine curves with a frequency ratio corresponding to a perfect third, using the colors of the Portuguese flag (green and red):



How to improve LLM Output?

- Use the latest LLMs.
- Experiment with different LLMs.
- Experiment with prompt engineering techniques.
- Use a RAG system.
- Use a CoT or MoE model.
- Combine LLMs with other tools.
- Fine-tune or distil an LLM.



Thanks for listening

Matthew.England@coventry.ac.uk

