GLOBAL FIRE SAFETY EXPERTS GROUP

AI in Engineering – Fundamental Understanding

Francis Galton - 1906

- Around 800 people were asked to estimate the weight of an Ox

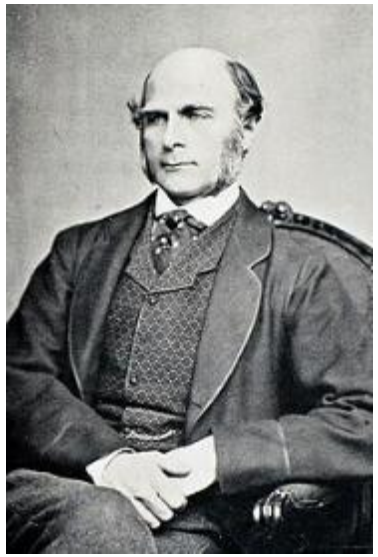

- Actual weight = 1,198lb
- Average guessed weight = 1,197lb

GLOBAL FIRE SAFETY EXPERTS GROUP

Oxen to Algorithms

- What Galton discovered accidentally in 1906 is, in essence, what modern AI does deliberately today.

- It aggregates billions of small, probabilistic ‘guesses’ to find a likely answer.
  In both cases, intelligence emerges not from any one mind, but from the ***aggregation of patterns***.

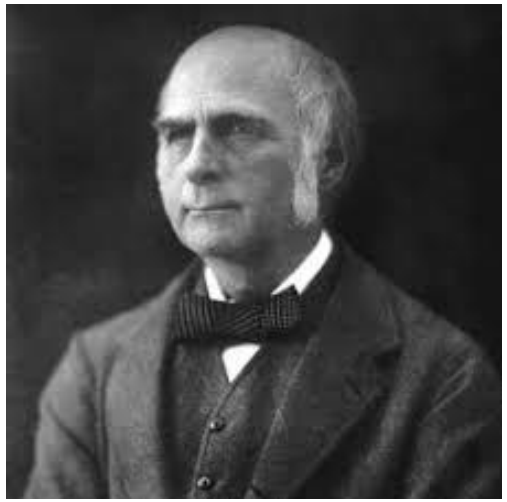

GLOBAL FIRE SAFETY EXPERTS GROUP

AI Building a Sentence

- The…

Next Word	Probability
cat	25%
man	20%
dog	15%
car	10%
biscuit	5%
velociraptor	0.5%
hydraulic	0.1%

GLOBAL FIRE SAFETY EXPERTS GROUP

AI Building a Sentence

- Sentence builds: **The cat…**

Next Word	Probability
sat	35%
ran	20%
sleeps	10%
exploded	0.4%

GLOBAL FIRE SAFETY EXPERTS GROUP

AI Building a Sentence

- **Full sentence:** ***“The cat sat on the mat.”***

- Eventually it builds a whole sentence, not because it understands what a cat or a mat is, but because those words often appear together.

The Prediction Space

- “The cat sat on the mat.”
  Probability ≈ **0.84%**

- The biscuit exploded on the keyboard.”
  Probability ≈ **0.00018%**

- Sometimes, the model wanders into strange territory.
  The second sentence is possible, just extremely unlikely.
  The AI doesn’t know which one makes sense, only which is more common.

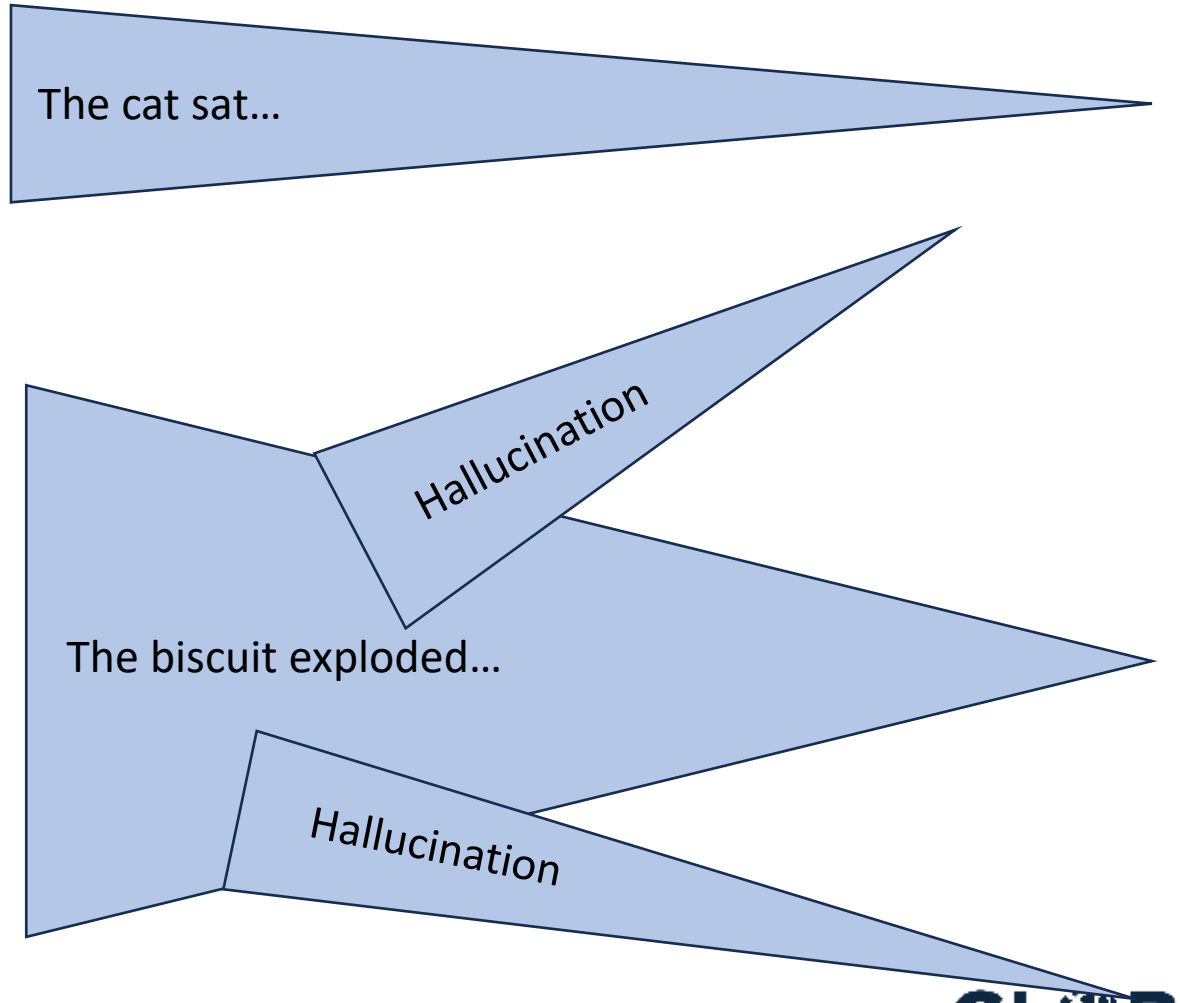

GLOBAL FIRE SAFETY EXPERTS GROUP

Hallucination

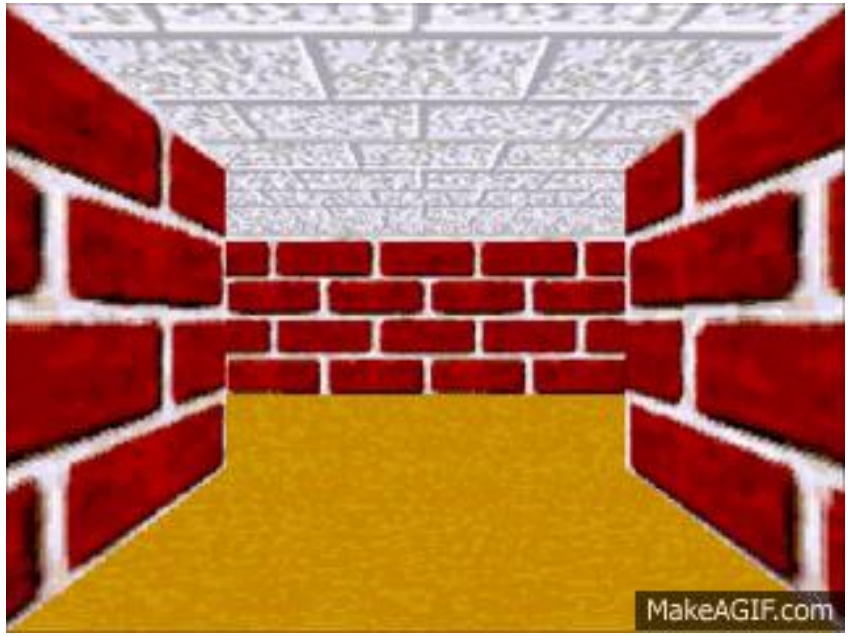

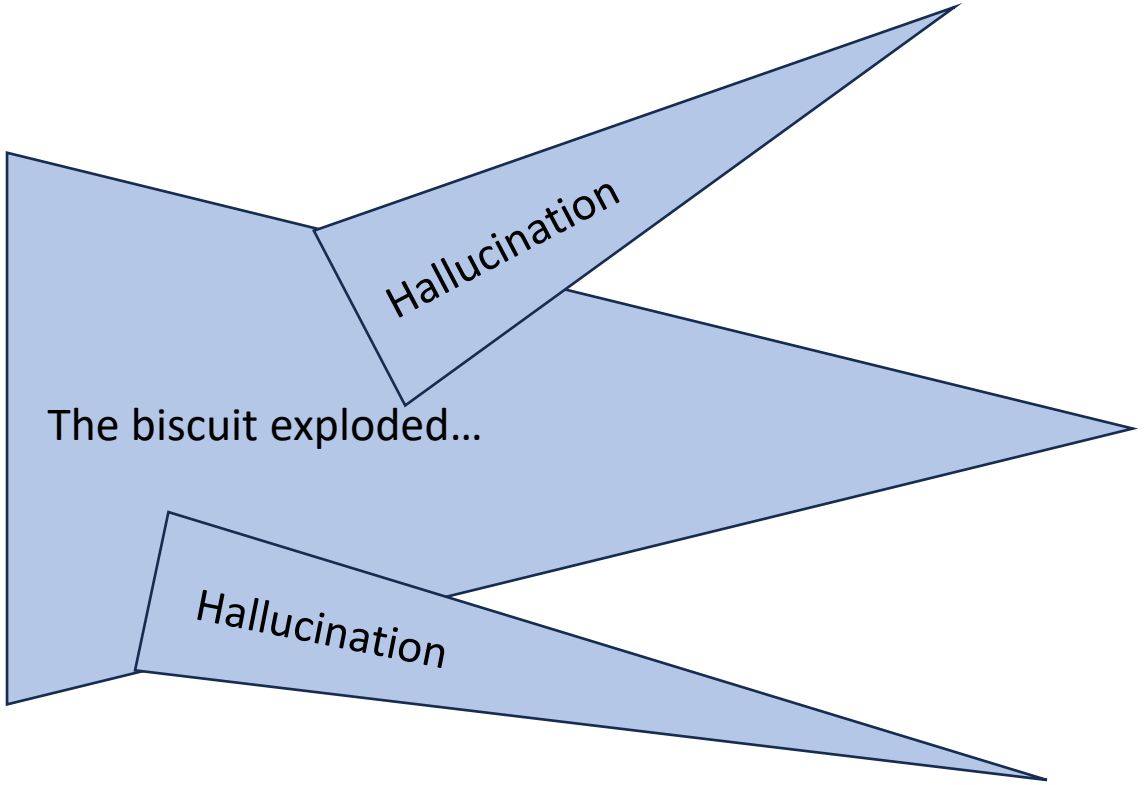

AI Guesses

- “The cat sat on the mat”

- Now imagine the same process, but with technical language.
  The model is still just guessing, and if most of its training data is American,
  it will guess *NFPA* long before *Approved Document B – Context is important!*

The Semantic Organisation: Ikea not Amazon

GLOBAL
FIRE SAFETY EXPERTS GROUP

The Showroom manager

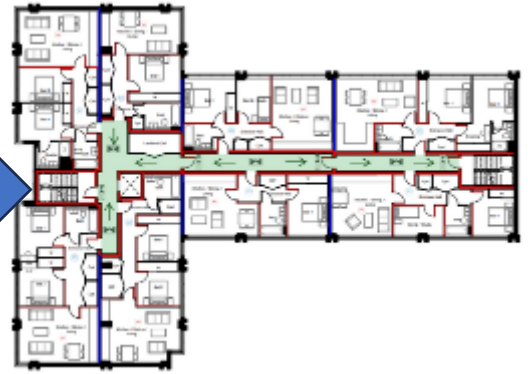

Cross-lingual token substitution

“Variable speed fans and control systems use pressure sensors at multiple (مستوياتtop, middle, bottom) to modulate airflow dynamically”

- ”مستويات“ mustawayāt - means ‘levels’ in Arabic

And:

“In those exposed हिस्सा, there are no windows, doors, or interfaces that people use…”

- हिस्सा means “areas” or “parts” in Hindi

What Happened?

- Let's assume we are back in IKEA and want a lamp.
- The model has walked into the right part of the IKEA showroom, the living room display, and picked a lovely standing lamp.
- It’s the correct type of object, in the correct context.
- BUT, the model has picked the lamp with a European plug, not a UK one. The idea is right, but the label isn’t.

Best Example vs Valid Example

Right pattern, wrong context: A child's brain

Pattern recognition,
just wrong context

GLOBAL
FIRE SAFETY EXPERTS GROUP

- An LLM (an AI) doesn’t understand, it ‘just’ pattern matches at scale

GLOBAL FIRE SAFETY EXPERTS GROUP

Context Anchoring

- Retrieval-Augmented Generation, or RAG, anchors the context – like a ship.
- The ship can still move - it can turn with the tide, and swing in the wind - but only within a safe and controlled radius.
- By anchoring the model to our own information; i.e. to verified standards, reports, and procedures; we give it freedom to reason, but not to wander (too far).

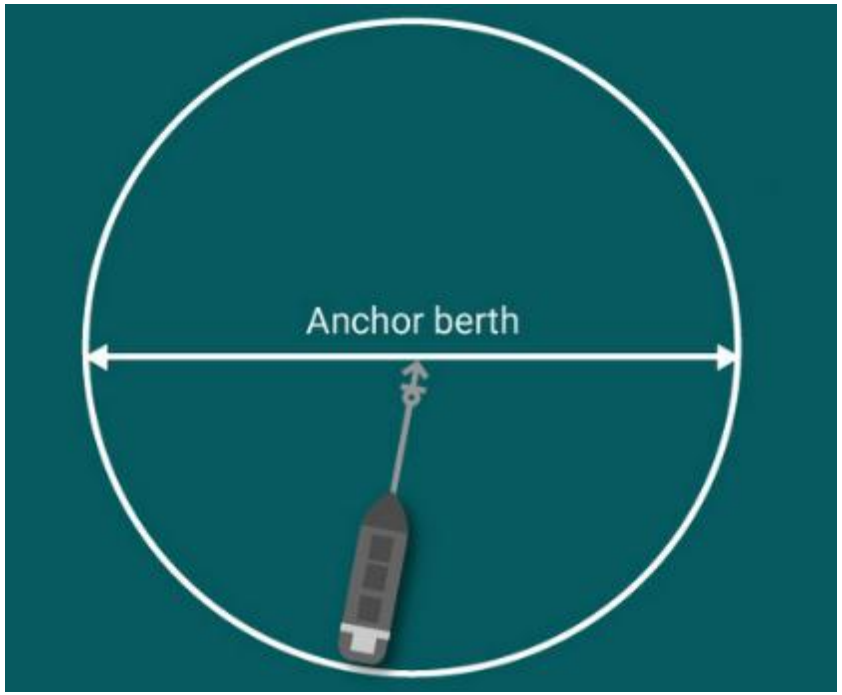

GLOBAL FIRE SAFETY EXPERTS GROUP

It’s a Tool

- Think of AI like a hammer. It’s not inherently safe or unsafe - **it depends on how it’s used**.
- In careful hands, with the right technique and purpose, it can produce finely crafted furniture.
- In another setting, it can demolish a wall; a different skill.
- But used carelessly, without training, safeguards, or supervision, that’s when it becomes dangerous.
- The tool didn’t change, only the context, the controls, and the user’s intent did.

GLOBAL FIRE SAFETY EXPERTS GROUP

Key Challenges in AI Adoption

- **Lack of transparency and control** — it is often unclear how AI models are trained, what data they use, or how decisions are reached.
- **Data ownership and privacy risks** — feeding sensitive documents into third-party AI tools raises serious questions about where data goes and who controls it.
- **Governance, bias, and regulatory compliance** — AI can generate biased or inaccurate outputs; organisations must ensure ethical use and compliance with GDPR and the EU AI Act.
- **Fragmented infrastructure** — legacy systems, siloed data, and inconsistent tooling slow down AI deployment and create technical debt.

GLOBAL FIRE SAFETY EXPERTS GROUP

How Organisations Engage with AI

- Most organisations don’t build AI from scratch. There are three practical routes:

- **Buy an off-the-shelf tool** — fast and low-cost to start, but you hand over control of your data, get limited transparency, and depend on a vendor’s decisions about training and updates.

- **Fine-tune an existing model on your own data** — the middle path. You adapt a foundation model to your domain — fire safety regulations, inspection records, approved documents — without building from the ground up. Better accuracy, more control.

- **Anchor with RAG** — as covered earlier, connecting any model to your verified documents gives domain accuracy without retraining. The lowest-risk route into professional use.

GLOBAL FIRE SAFETY EXPERTS GROUP

What This Means for Your Organisation

- **Tie AI to specific outcomes** — don't adopt AI for its own sake. Identify where it saves time, reduces error, or improves compliance accuracy in your existing workflows.
- **Start with low-risk pilots** — test AI on internal document search or report summarisation before applying it to live compliance decisions.
- **Establish governance before you deploy** — define who reviews AI outputs, how errors are caught, and what the escalation path looks like. Policy first, technology second.
- **Move early, but move carefully** — early adopters in safety-critical sectors will shape industry standards. Getting this right now builds competitive advantage and professional credibility.

GLOBAL FIRE SAFETY EXPERTS GROUP

Global Fire Safety Group and BridgeAI

- Earlier this year I met with Keith from Global Fire Safety Group to explore how the business could begin adopting AI responsibly and in a way that genuinely supports the work the team already does.

- Through the **BridgeAI programme,** backed by Innovate UK and the Alan Turing Institute , Global independent expert advice, and a structured pathway to adopt AI safely, without needing to figure it out alone.

- The conversation started with the same questions this presentation has been asking: **what can AI actually do, where does it fail, and how do you deploy it in a way you can defend?**

Human in the Loop

- We have established that AI is a pattern-matcher, not a reasoner. It can hallucinate, pick the wrong context, and confuse a valid answer with the best answer.

- In fire safety, where a wrong answer can cost lives, **human judgement remains the ultimate safety system.**

- AI can draft, summarise, flag, and suggest, but **domain experts decide** what is safe, compliant, and legally defensible. That accountability cannot be delegated to a model.

- Every AI workflow in a professional context should have a defined point at which a qualified person reviews, approves, or overrides the output. That is not a limitation of AI, it is responsible use of it.

GLOBAL FIRE SAFETY EXPERTS GROUP

ISO 42001: A Framework for Safe AI Adoption

- ISO 42001 is the international standard for AI management systems. Think of it as the ISO 9001 of AI. It gives organisations a structured, auditable way to deploy AI responsibly.

- **Risk management and transparency** — document what the AI does, where it is used, and what the failure modes are.

- **Accountability and governance** — assign clear responsibility for AI outputs, with defined review and escalation processes.

- **Familiar territory for fire safety professionals** — the framework mirrors other safety-critical management standards. If you already work within a quality or safety management system, ISO 42001 follows the same logic.

GLOBAL FIRE SAFETY EXPERTS GROUP

Next Steps for AI Adoption

- **1. Identify one or two high-value use cases** — where does AI save the most time or reduce the most risk in your current workflows?
- **2. Launch a low-risk pilot** — a document assistant, report summariser, or internal search tool. Build confidence before touching live compliance decisions.
- **3. Establish governance and ethical use policies** — aligned to ISO 42001. Define review processes, accountability, and escalation paths before you deploy.
- **4. Monitor, iterate, and improve** — collect feedback from your domain experts, retrain or refine the model, and expand scope only when trust is established.

GLOBAL FIRE SAFETY EXPERTS GROUP