MINNESOTA STATE WORKGROUP ON AI FOR HEALTHCARE

AI IN PLAIN ENGLISH

JANUARY 26, 2026

AI in Health Care is HOT

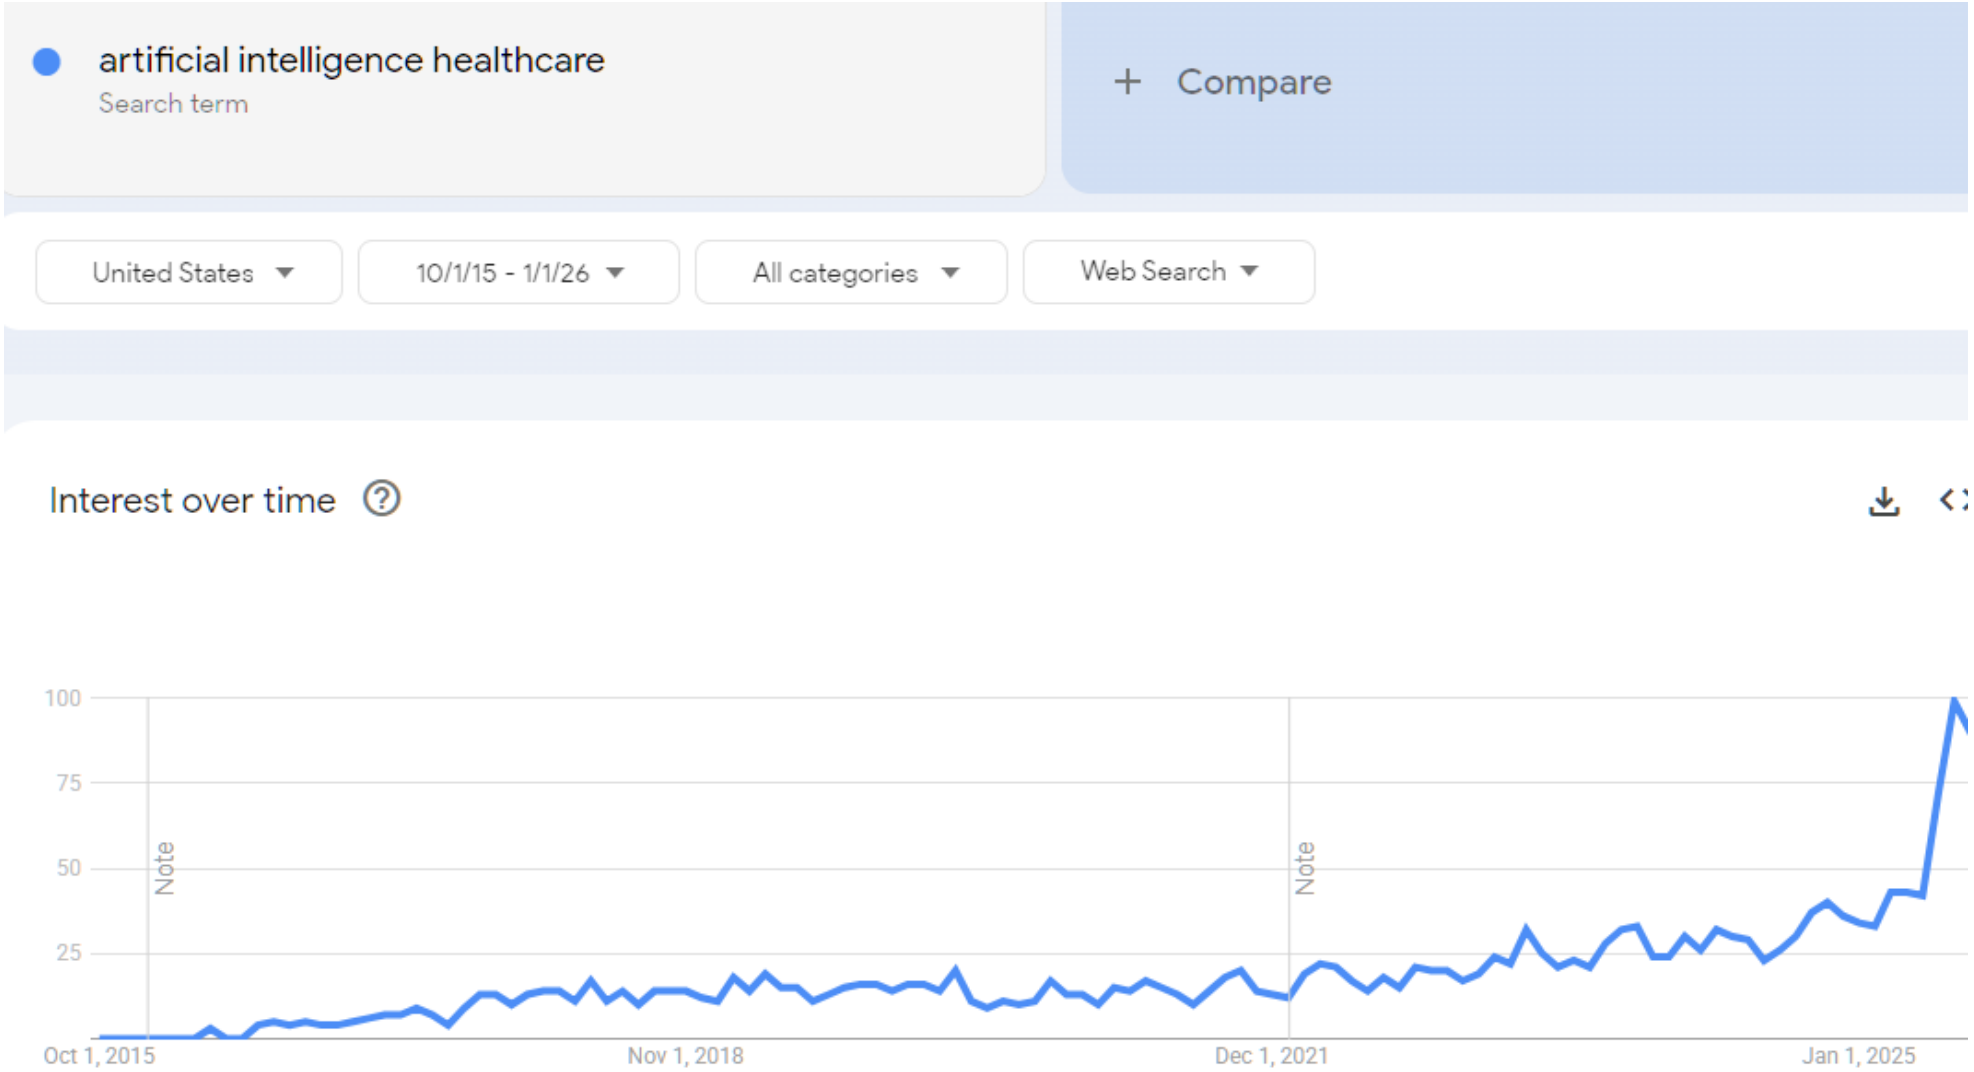

NVIDIA and Lilly Announce Co-Innovation AI Lab to Reinvent Drug Discovery in the Age of AI

Companies to Jointly Invest up to $1 Billion Over Five Years in Infrastructure and Research

January 12, 2026

Chai Discovery Announces Collaboration with Eli Lilly and Company to Accelerate Biologics Discovery

January 12, 2026 | 2 min read

PR Newswire

CytoReason extends its collaboration with Sanofi to advance AI-driven drug discovery

NEWS PROVIDED BY
CytoReason →
Jan 09, 2026, 07:00 ET

MedCityNews

STARTUPS, HEALTH TECH, ARTIFICIAL INTELLIGENCE, PROVIDERS

Anthropic Follows OpenAI Into Healthcare: How Do Their Platforms Compare?

AI rivals Anthropic and OpenAI are both expanding their large language models into healthcare. Anthropic is blending its enterprise and consumer tools in a single platform, while OpenAI is separating its consumer-facing ChatGPT Health from its industry-focused OpenAI for Healthcare. They are both targeting patients, providers and researchers with AI tools for tasks ranging from personal health insights to coding to prior authorization.

By Katie Adams on January 11, 2026 7:39 pm

October 16, 2025 | Company News

Microsoft extends AI advancements in Dragon Copilot to nurses and partners to enhance patient care

UNITEDHEALTH GROUP

Ensuring responsible use of AI in health care

Discover how we responsibly use artificial intelligence solutions to help people live healthier lives and make the health system work better for everyone.

MAYO CLINIC | News Network

MEDICAL INNOVATION

GE HealthCare and Mayo Clinic unveil GEMINI-RT, a bold research collaboration in radiation therapy and advanced cancer care

Tom Millikin
December 3, 2025

Health systems begin activating Epic's new features

…but trying to make sense of the hype cycle can be overwhelming

90+ Healthcare AI Startups To Watch

Imaging & Diagnostics

ablacon, aidence, aidoc, Healthy.io, PathAI, RadiomicsAI, Aillis, EXO, 深睿医疗 deepwise, 推想科技 infervision, sight, enlitic, EXAMWIZARDS, freenome, JLK INSPECTION, zebra, BUTTERFLY Network, Inc., IBEX, IDx, SIGTUPLE, 数坤科技, ULTROMICS, PAIGE, VUNO, VoxelCloud, Lunit, VIZ

Drug Discovery

Accutar Biotech, BenevolentAI, DIP 深度智耀, FINCH THERAPEUTICS, APEIRON, BlackThorn THERAPEUTICS, Frontier MEDICINES, insitro, Atomwise Better medicines faster., Exscientia DRIVEN BY KNOWLEDGE, engine BIOSCIENCES, JUVENESCENCE, Notable Labs, OWKIN, RECURSION pharmaceuticals, 药研社, Standigm, XtalPi, twoXAR, SPRING DISCOVERY

Predictive Analytics & Risk Scoring

BenchSci, lu, 森亿智能 SYNYI·AI, saama, GNS HEALTHCARE, Medial EarlySign, TEMPUS, ikioo, medopad, Welltok, nference

Genomics

DECIPHER BIOSCIENCES, SOPHIA GENETICS, Syntekabio, WuXiNextCODE

Fitness

FINC, Vi

Virtual Assistant

babylon, Curai, mfine, notable, buoy, K, Suki

Hospital Decision Support

binaryfountain, Cardinal Analytx Solutions, Qventus, gauss, Olive, OM1, Ken Sci, NextHealth TECHNOLOGIES, LeanTaaS, Medical INFORMATICS

Clinical Trials

deep, HEALTH[at]SCALE TECHNOLOGIES, teckro.

Nutrition

NURITAS LIFE-CHANGING INGREDIENTS, WHOLE BIOME

Remote Monitoring

biofourmis, ContinUse Biometrics, DENTAL MONITORING Connected orthodontics

Compliance

DATAVANT, PROTENUS

Mental Health

mindstrong

Created by You. Powered by CBINSIGHTS

HEALTHCARE AI MARKET MAP

SAPPHIRE VENTURES

Voice AI

Elise AI
Hippocratic AI — Do No Harm —
Assort Health
hellopatient
hyro
ELLIPSIS HEALTH
INFINITUS
Clarion

Clinical Co-Pilot

ABRIDGE
OpenEvidence
Ambience
regard
commure
heidi
Freed
sully.ai
Nabla
Suki
Tandem

Image Analysis & Diagnostics

Rad AI
HOPPR
iz.ai
New Lantern
PathAI
SIRONA
aidoc
cleerly

Home Health

Verse Medical
Roger.
EnzoHealth
Andy

AI-Chart Review

LAYER HEALTH
Charta
Pharos
Brellium

Record Retrieval

predoc
Codes Health
zus
Credo

Revenue Cycle Management (RCM)

Intake, Eligibility & Prior Auth

tandem
Mandolin
INSIGHT HEALTH
Tennr
latent
Plenful
Anterior
SILNA
valerie health
EXACARE

AI-Coding

CODAMETRIX
nym
FATHOM
Parallel
AKASA

AR, Claims & Billing

candidhealth
stedi
adonis
Camber

Denials and Underpayment

Claim Health
aegis
Cofactor

Patient Billing

Collectly
cedar

Behavioral Health

slingshot AI
limbic
eleos
Alma
growtherapy
Headway

Back-Office & Operations

Clarium
trek health
Nitra
notable
translucent
keragon
BitBoard

AI EHRs

prompt
Canvas:
Healthie
UrgentIQ
Jane
ARYAHealth
Simplify Your Practice
Elation

Sapphire has internally assigned categories to the companies described above. Although Sapphire believes that such categorizations are reasonable, they are subjective in nature and may be categorized differently by other market participants.

Sapphire Ventures investment

My promise for today

Plain English Fundamentals

Practical Tips & Tricks

Zero BS

AI/ML CONCEPTS

01

THE BASICS – AI, ML, NEURAL NETWORKS

Artificial Intelligence

Engineered systems that can perceive the world, ingest data, reason about that data and take meaningful action.

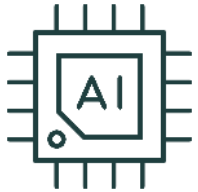

The term was originally coined in 1956 at Dartmouth and researchers then predicted near human intelligence (GAI) within a decade.

Models can then use the learned patterns to predict relationships and process previously unknown data.

Machine Learning

Mathematical and statistical models (algorithms) that perform analysis on large quantities of historical data and identify patterns.

Neural networks

A neural network is a type of machine learning algorithm that uses a structure inspired by interconnected neurons in the human brain.

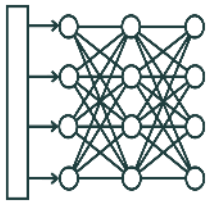

Neural networks that use many layers of these interconnected nodes (“neurons”) are called “deep learning” algorithms

THE BASICS – GENERATIVE AI, LLMS

Generative AI

Systems that create new content – text, images, data, audio, or video – based on patterns learned from existing data.

Generative AI entered the public consciousness with the arrival of Dall-E for image generation in 2021 and ChatGPT for text generation in November of 2022.

They are based on patterns learned from vast amounts of training data. Today's LLMs are based on the Transformer architecture first described in 2017.

Large Language Models (LLMs) / Generative Pretrained Transformers (GPTs)

A subset of generative AI, Large Language Models process and generate human-like text by predicting the next word or phrase one at a time.

Open AI's ChatGPT, Anthropic's Claude, Meta's LLama, Google's Gemini, etc. are examples of open-sourced and closed-style LLMs developed by leading enterprises. There are hundreds of other open-source LLM models available today.

THE BASICS - TOKENS

- Tokens are text chunks (like words/parts of words) that AI uses to process language
- Using tokens is efficient - common patterns get reused while rare words split into pieces, saving compute

Bjorn Austraat – CEO and Founder of Kinetic Cognition | AI Innovator & Practitioner

Bjorn Austraat is an award-winning and industry-recognized leader in Artificial Intelligence (AI) and innovation, with over two decades of extensive experience across high-tech, financial services, and telecommunications sectors. With a robust track record of driving transformative changes, Bjorn has successfully led initiatives in global businesses ranging from $10M to $95B, applying a continuous improvement approach to set and

```
Bjorn Austraat - CEO and Founder of Kinetic Cognition | AI Innovator &
 Practitioner
Bjorn Austraat is an award-winning and industry-recognized leader in
 Artificial Intelligence (AI) and innovation, with over two decades of
 extensive experience across high-tech, financial services, and
 telecommunications sectors. With a robust track record of driving
 transformative changes, Bjorn has successfully led initiatives in global
 businesses ranging from $10M to $95B, applying a continuous improvement
 approach to set and achieve strategic enterprise goals.
```

THE BASICS – LLMS AND TOKENS

At their most basic level, LLMs are “Next Best Token” prediction machines

- User supplies a piece of text (the prompt)
- LLMs chunk the prompt into tokens and make sense of all tokens in context
- LLMs then answer the question: given this sequence of tokens, what is the next best token?
- Rinse, repeat until a stop token is issued

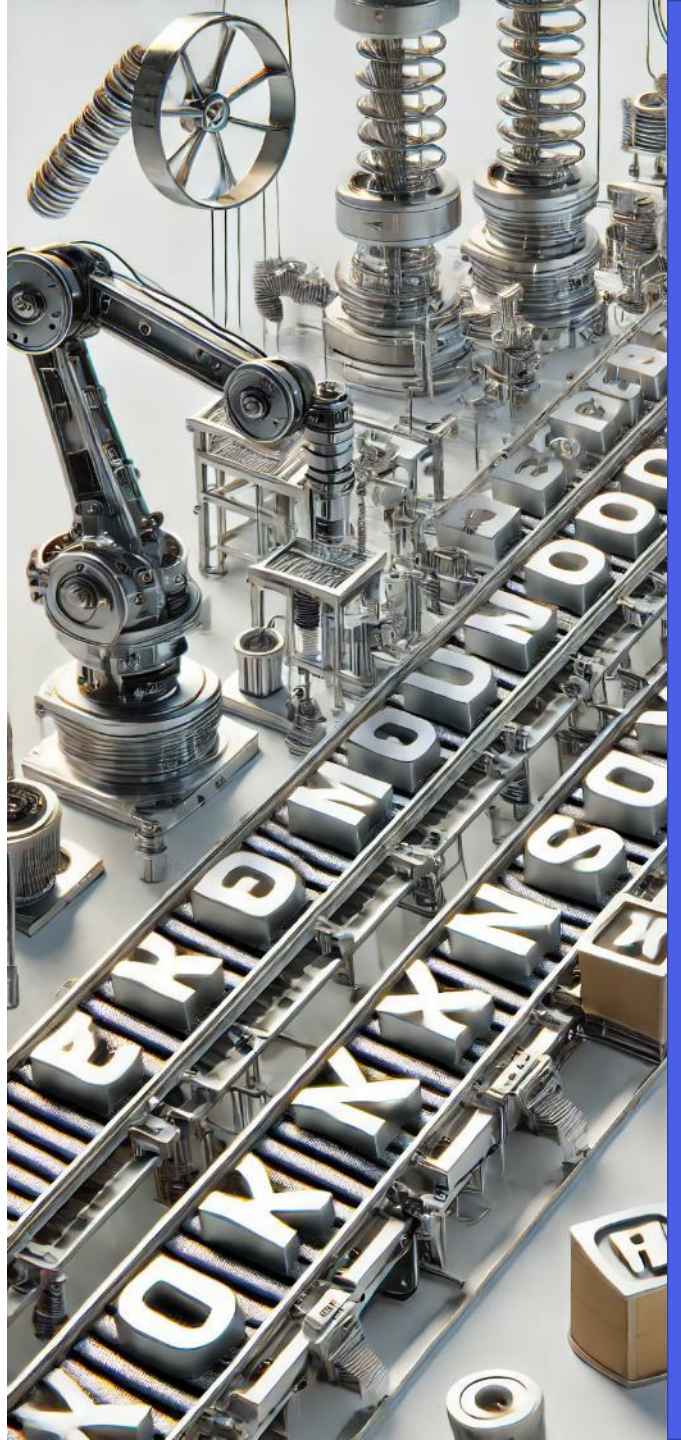

VECTORS & VECTORIZATION

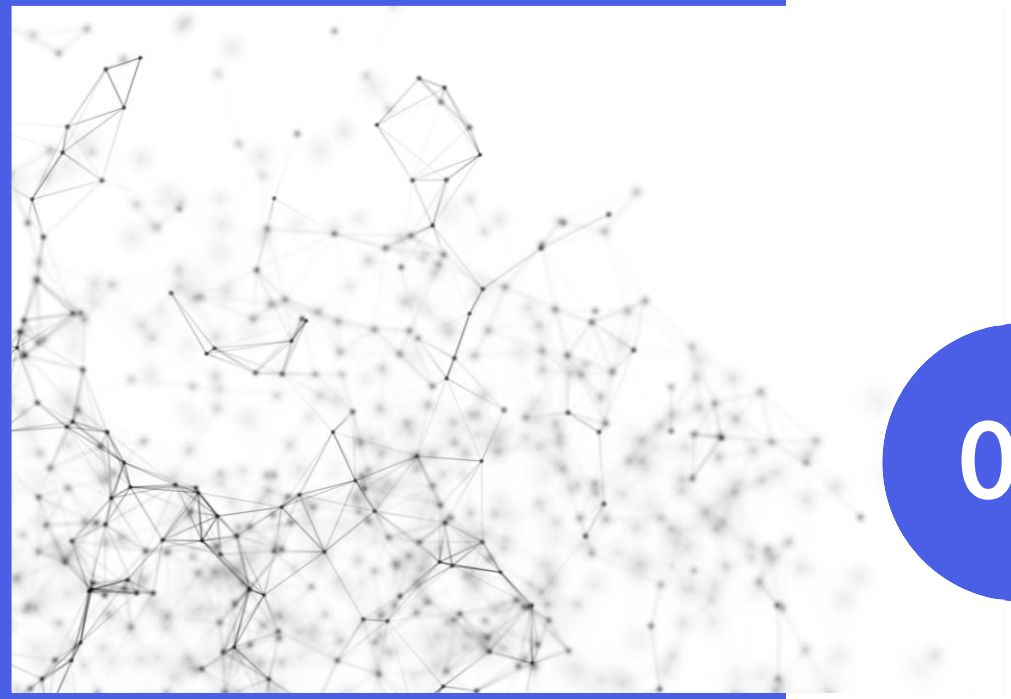

02

TO HELP COMPUTERS MAKE SENSE OF **WORDS**

WE TURN THEM INTO **NUMBERS**

"features"	Age	Authority	Shoulder_Width	"vector"
boy	10	10	10	[10,10,10]
girl	10	10	7	[10,10,7]
king	75	90	70	[75,90,70]
queen	75	90	60	[75,90,60]
prince	25	50	45	[25,50,45]
princess	25	50	35	[25,50,35]
man	50	40	50	[50,40,50]
woman	50	40	40	[50,40,40]

NOW THAT WE HAVE TURNED WORDS INTO MEANINGFUL VECTORS, WE CAN SHOW HOW WORDS RELATE TO EACH OTHER (IN SPACE)

THIS IS CALLED **EMBEDDING**

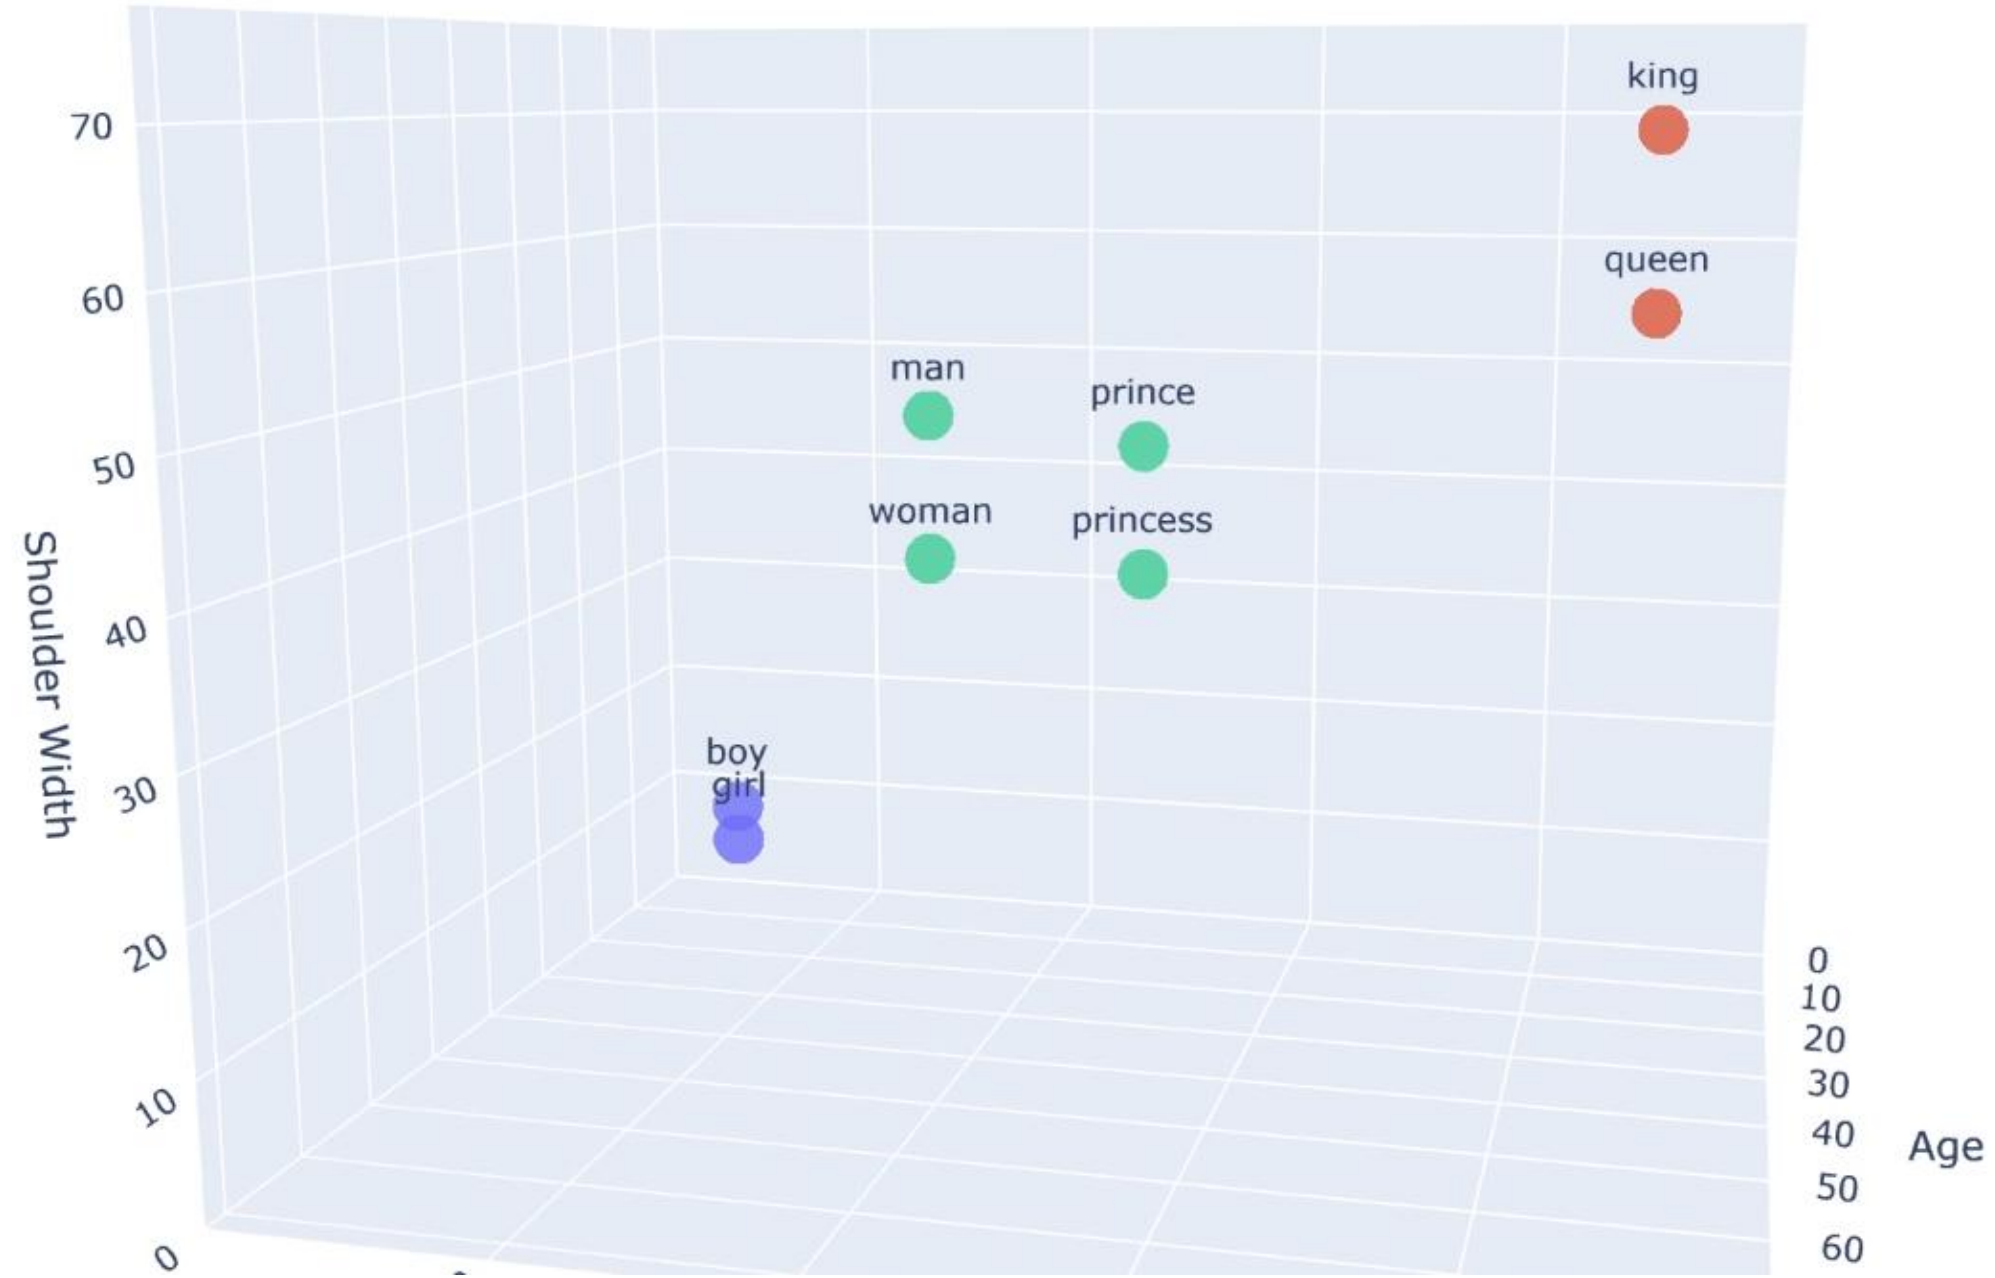
king
queen
man
prince
woman
princess
boy
girl
70
60
50
40
30
20
10
0
Shoulder Width
0
10
20
30
40
50
60
Age

Vectors let us express relationships between words using math

	Age	Authority	Shoulder
man	75	90	50
minus	0	0	10
= woman	75	90	40

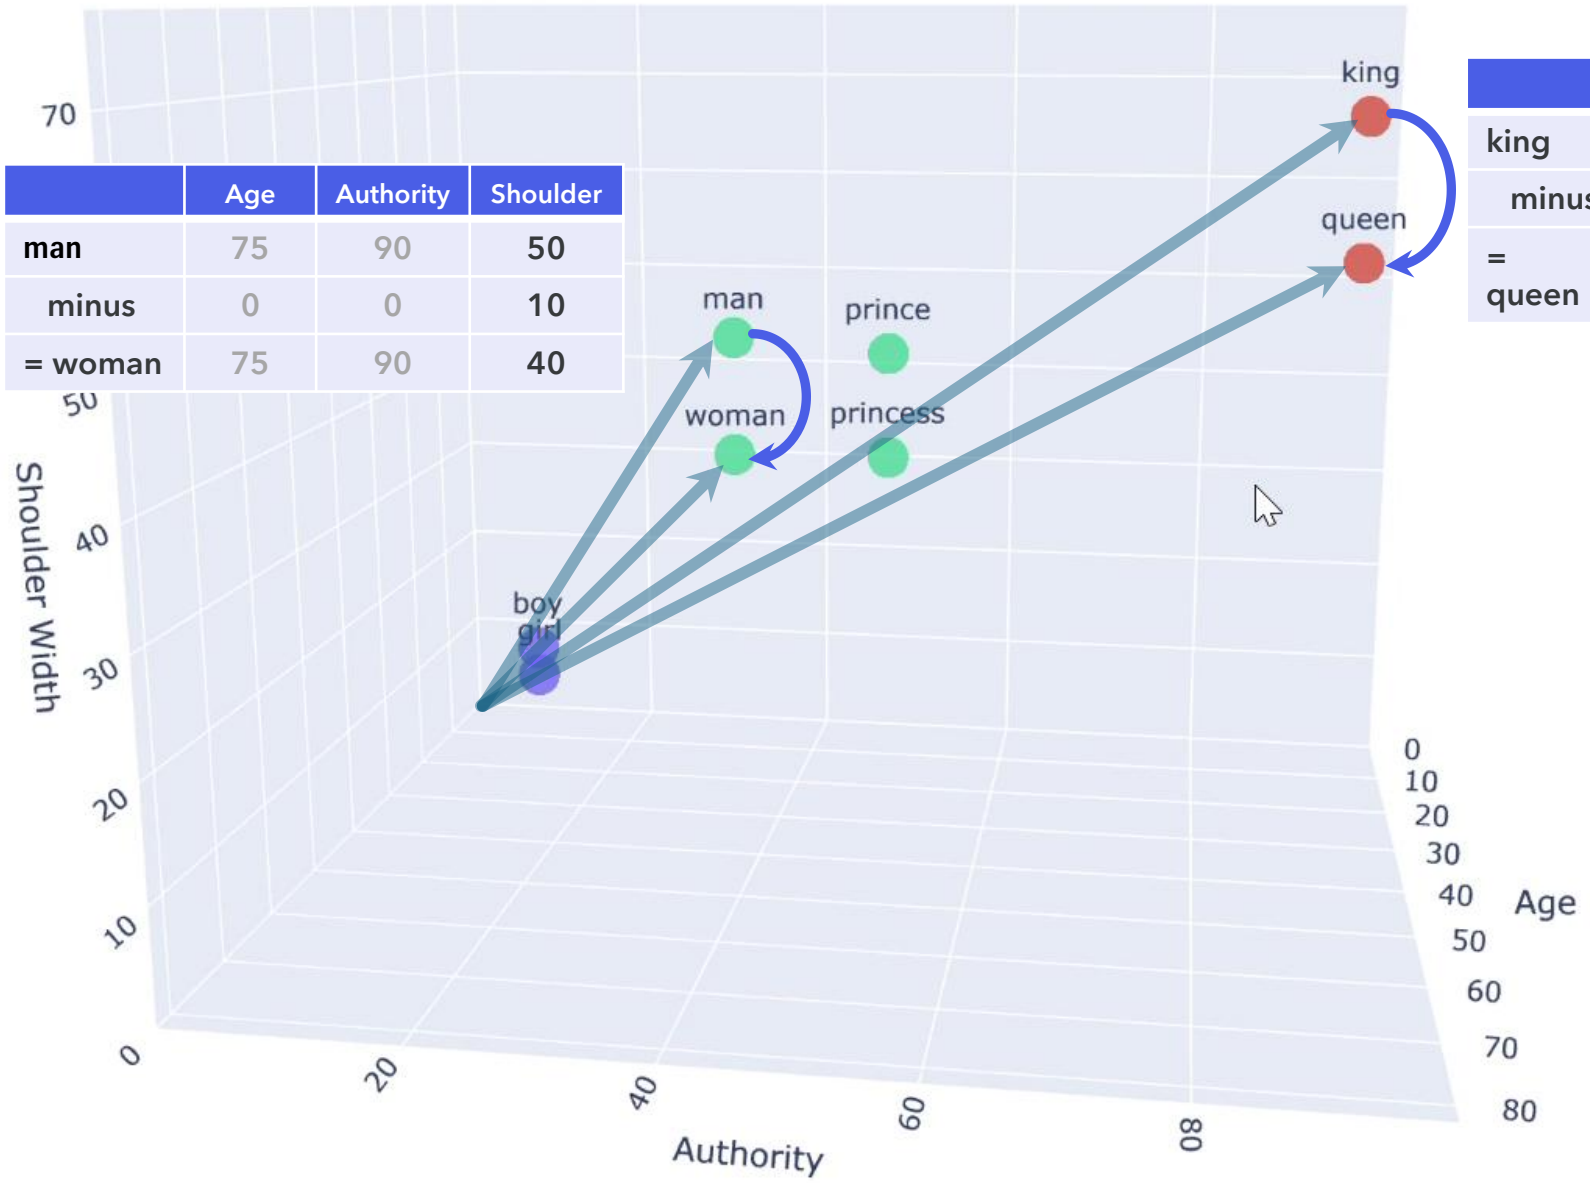

	Age	Authority	Shoulder
king	75	90	70
minus	0	0	10
= queen	75	90	60

Sophisticated embedding approaches use hundreds or even thousands of dimensions

These dimensions can’t easily be translated into simple concepts like age or shoulder width but represent an abstract way of seeing the world

For example, OpenAI provides an embedding model that maps text to

1,536 dimensions

word	embedding_0	embedding_1	embedding_2	embedding_3	embedding_4	embedding_5	embedding_6
boy	0.041855	-0.01652	-0.0266	0.063359	0.006049	-0.06054	0.010412
girl	0.032278	-0.02934	-0.02708	0.024227	0.021979	-0.05546	0.032454
king	0.037222	-0.02209	0.051943	0.000141	-0.01373	-0.02303	-0.02378
queen	0.043817	-0.03984	0.044738	0.021715	-0.00485	-0.01212	-0.0125
prince	-0.00776	-0.00536	0.024468	-0.00185	-0.02078	-0.01486	-0.01833
princess	0.038494	-0.01822	-0.00088	0.017732	-0.02555	0.002838	-0.01413
man	0.038761	0.000384	0.016608	0.023913	-0.0094	-0.05897	-0.00966
woman	0.061607	-0.03538	-0.0189	0.070923	-0.00052	-0.04178	-0.00297

Features / Embedding dimensions

Heat map visualization of features

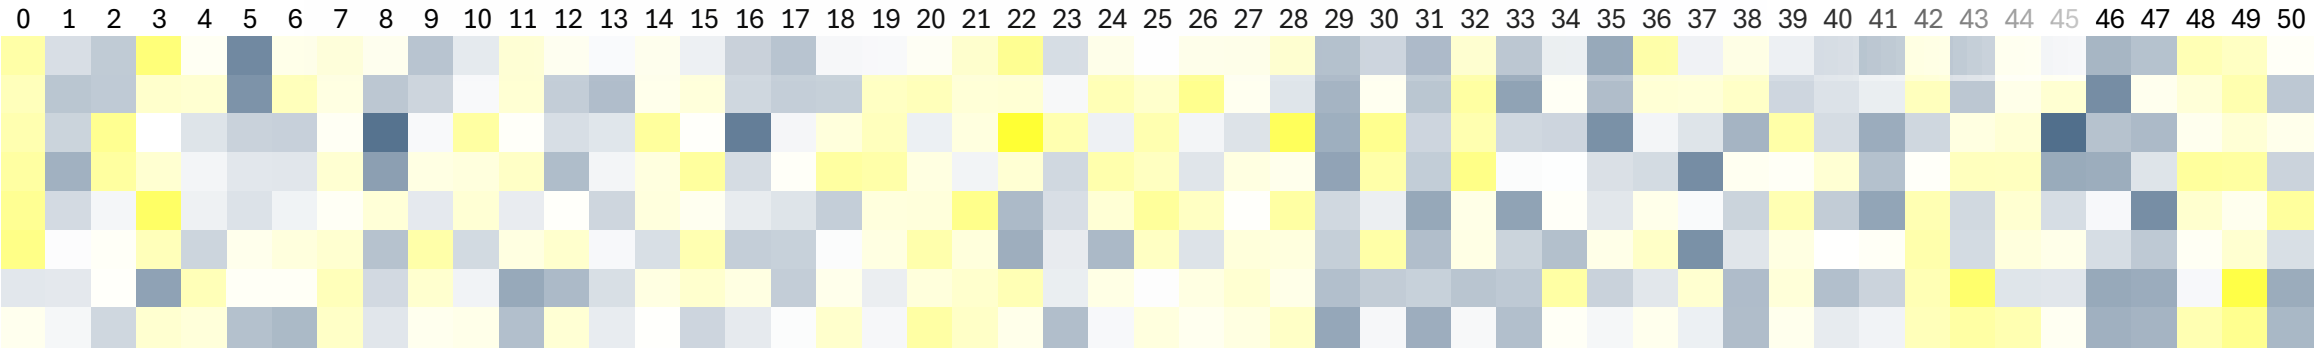

boy

girl

king

queen

dog

bird

machine learning

car

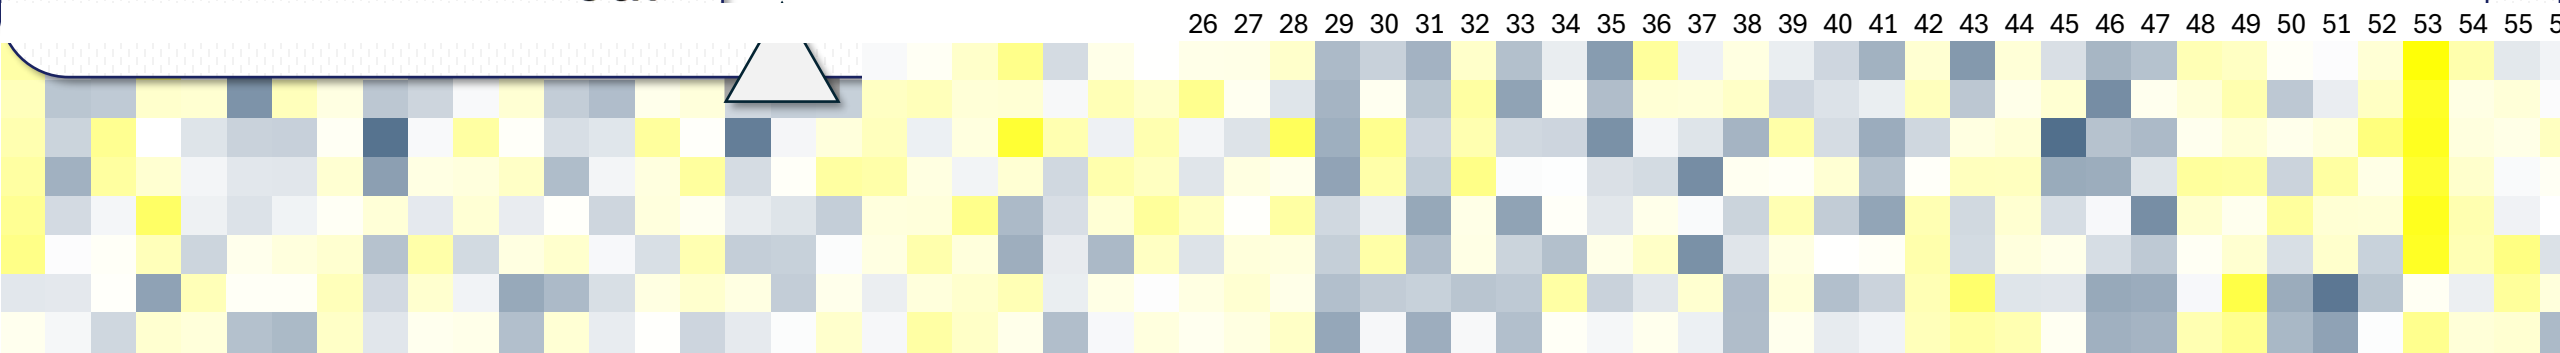

PRACTICAL TIPS

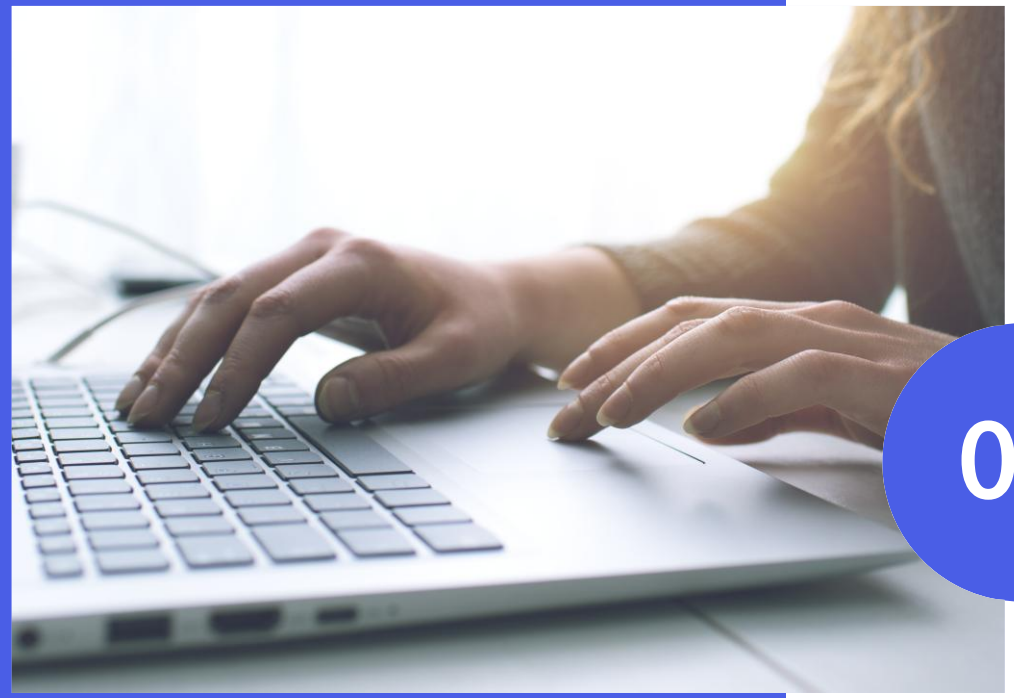

03

VENDOR CHECKLIST CHEAT SHEET

- **Can you switch LLM vendors if needed**
  Are you locked into one AI provider, or can you change without breaking everything.
- **Is our data kept private**
  Does anything we put into the system get reused to train someone else's AI.
- **Does the AI change over time**
  How do you know when the AI's behavior drifts or gets worse.
- **Can you explain its answers**
  Can you show how and why the AI gave a result.
- **How do you check quality**
  How do you make sure the answers stay accurate and useful.
- **What happens when it makes things up**
  How do you detect and fix false or misleading answers.
- **How secure is it**
  How do you protect the system from hacking, tricking, or data leaks.
- **Are there hidden vendors inside**
  Are other AI systems being used behind the scenes.
- **Who is responsible if something goes wrong**
  Who pays if there are legal, financial, or copyright problems.
- **Are the people running it verified**
  How do you prevent fake or fraudulent users and employees from getting access.

AI STRATEGY IN A BOX – 5 PILLARS

Platform

The foundation that lets AI exist beyond experiments. This is about having clean, accessible data and tools that teams can actually use.
Best practice: Build once, reuse often. Make data easy to find, models easy to run, and systems reliable enough that people trust them.

Governance

The rules that keep AI safe, fair, and out of trouble. This is about trust, not bureaucracy.
Best practice: Design controls early, make responsibilities clear, and monitor systems continuously instead of reacting after something breaks.

Talent

The humans who build, run, and rely on AI. This includes specialists and non-specialists alike.
Best practice: Invest in learning at all levels, reward collaboration, and document work so knowledge doesn't walk out the door

Communication

How people understand what AI is doing and why. If no one can explain it, support disappears quickly.
Best practice: Use simple stories, visible metrics, and repeat yourself more than feels necessary. Clarity beats cleverness**.**

Solution Portfolio

Deciding where AI is worth the effort. Some ideas save money fast, others change how the business works over time. You need both.
Best practice: Start with real problems, not cool technology. Be explicit about expected value and stop projects early if they don't pan out.

Visit the AI Nexus for more

One more thing…

AI moved from **Technology Topic** to **Leadership Competency**

"Can AI do this" is now **"Should we let AI** do this**?"**

AI Decisions **Are Easy to Start & Hard to Undo**

If you are responsible for decisions that compound over time, let's talk.

THANK YOU

LinkedIn: linkedin.com/in/bjornaustraat

Kinetic Cognition: kineticcognition.com

AI Nexus: bjornaustraat.substack.com