SQL
KONFERENZ
2026

Building Data Literacy

In the Microsoft Data Platform Era



In this presentation…

- **Explore the basics of Data Literacy**
  - How does it apply to everyone?
  - What makes it so important?
- **Find ways to measure Data Literacy**
  - In a Microsoft Environment
- **Discuss how to meaningfully improve Data Literacy**
  - What does ‘training’ mean?
  - Whose job is it even?
- **How AI influences all of this…**



SQL KONFERENZ 2026

smart
casual
datadesign



Oliver Engels

Oliver Engels **is CEO of oh22data AG**, a Microsoft **Gold Partner** in Germany specialized in **CRM and BI**.

His **interests** are on Azure, Data Governance and Integration, Machine Learning and visualization Tools like SSRS, Power BI and Tableau.

Oliver has worked with SQL Server **since version 6.5** and is a founding member and current **president** of German Microsoft Data Platform Community

smart
casual
datadesign



Katharina Ricken

Tries her best.

Only recently changed careers, going from **3D to data engineering**, and is very humbled to speak here.

What is Data Literacy?

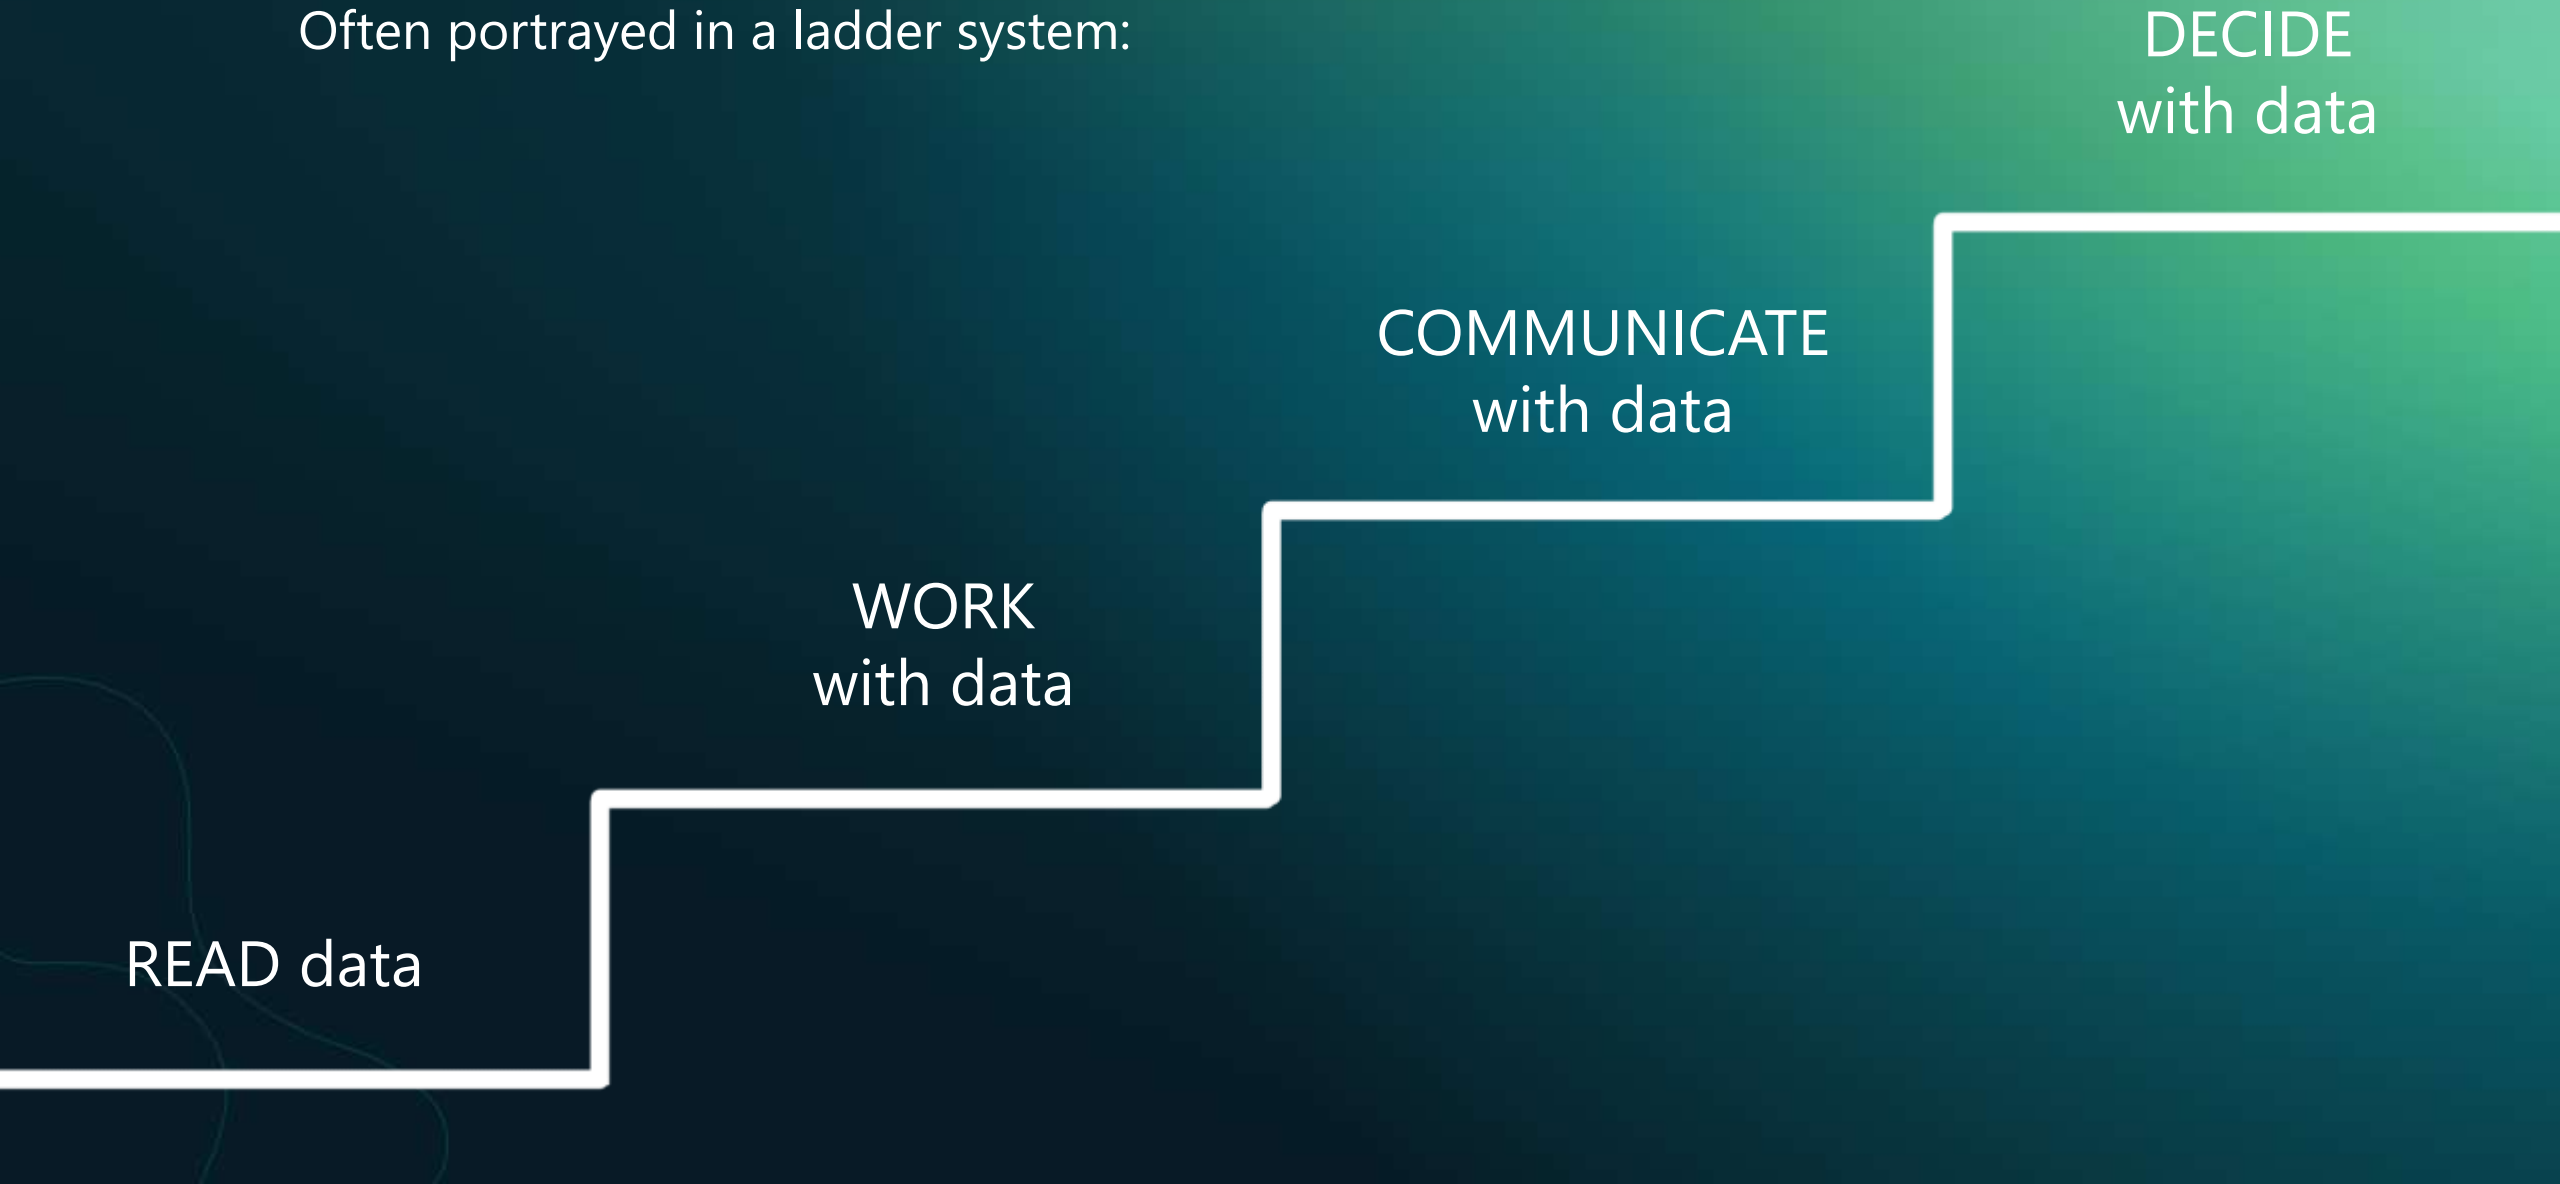
"The ability to explore, understand, and communicate with data in a meaningful way."
- Tableau

"Data literacy is the ability to read, write and communicate data in context,..."
- Gartner

"Data Literacy refers to the essential 21st century skill of collecting, managing, evaluating, and applying data in a critical manner."
- Wikipedia

What is Data Literacy?

Often portrayed in a ladder system:



DECIDE with data	=	Interpret results and act	→	Innovate competitively
COMMUNICATE with data	=	Present findings, argue insights	→	Influence
WORK with data	=	Manipulate, analyze, transform	→	Self-service
READ data	=	Understand charts, dashboards, datasets	→	Trust; Confidence

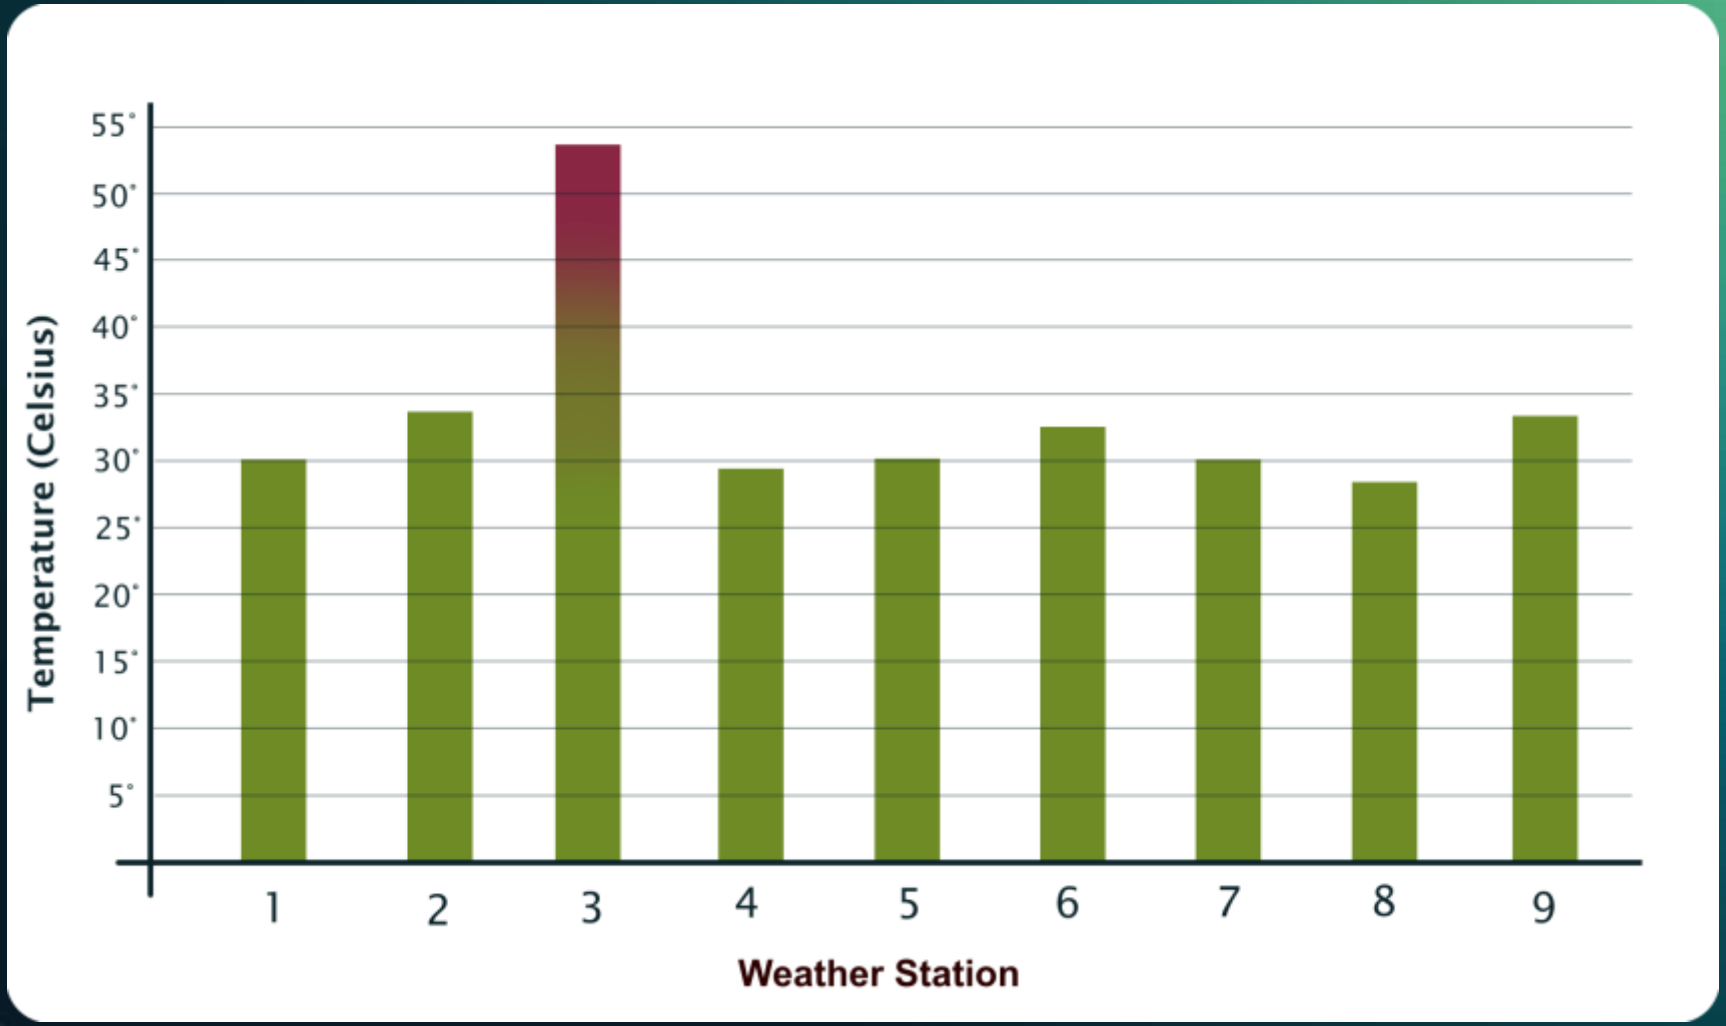
Why does Data Literacy matter?

In a self-service environment like **Fabric/ Power BI**, data literacy empowers users to explore and analyze data **without dependence** on IT

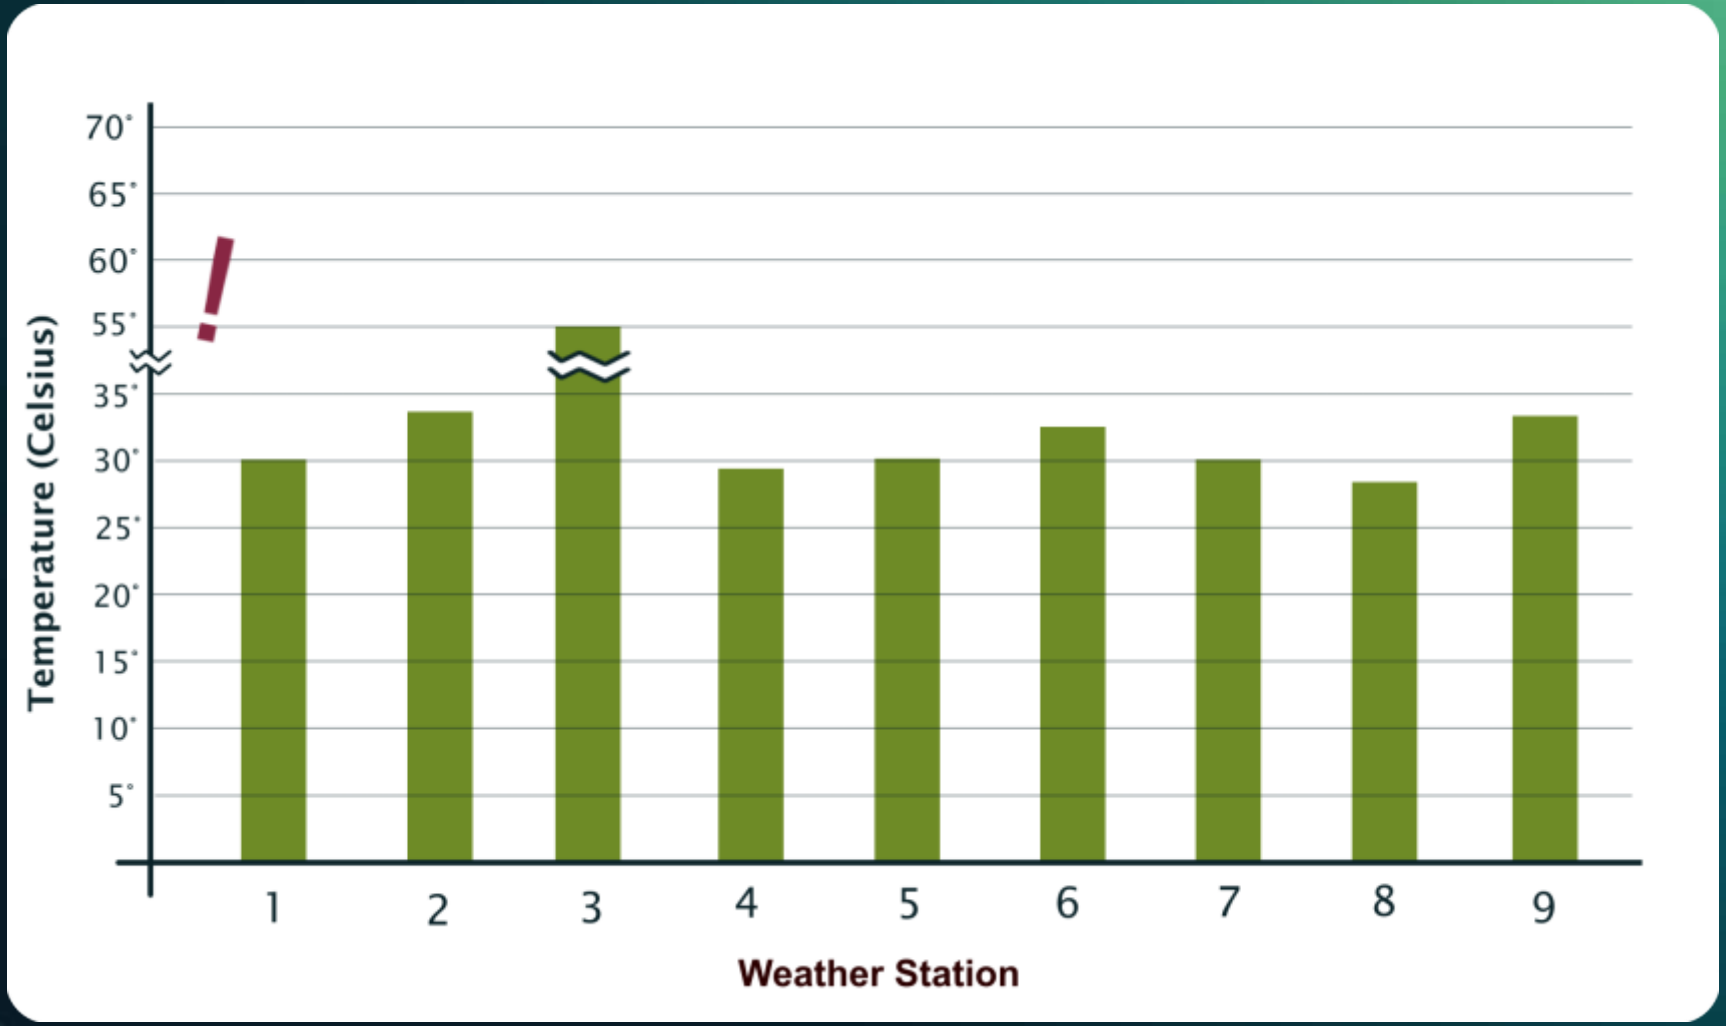
And without it?

- Poor decisions due to misinterpretation
- Mistrust or misuse of reports
- Adoption of new platforms stagnates
- Data governance suffers

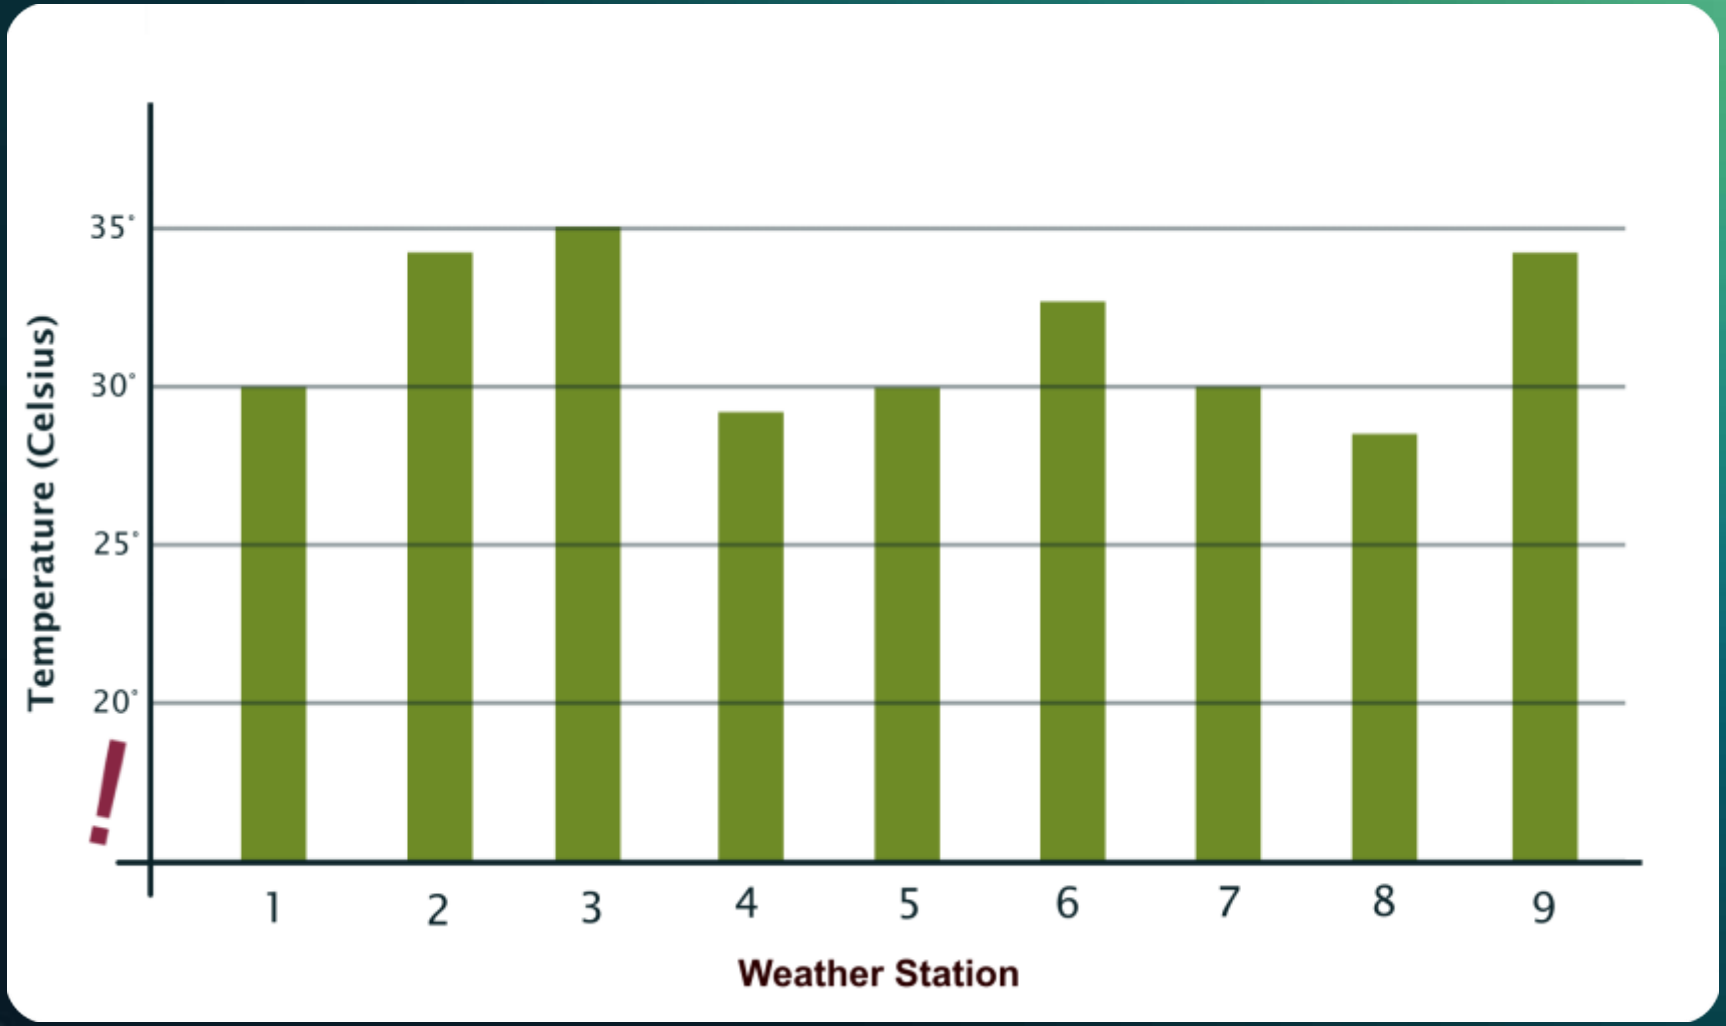
In a **day-to-day**, data-driven environment, data literacy empowers us to explore and analyze data **by depending** on our critical thinking skills.



In a **day-to-day**, data-driven environment, data literacy empowers us to explore and analyze data **by depending** on our critical thinking skills.



In a **day-to-day**, data-driven environment, data literacy empowers us to explore and analyze data **by depending** on our critical thinking skills.



In a **day-to-day**, data-driven environment, data literacy empowers us to explore and analyze data **by depending** on our critical thinking skills.



Sam Antic the Model

Saved my a** letting me use their toilet for free.

10/10 burger to go was good too



Carrie L'Atieun

Best burgers in town, just wish they'd do deliveries



Phil Terrows

Nice little establishment with a lot of charme.
Got order wrong once in countless visits and it is sometimes really full for such a small place, but I guess that's just proof how good it is :) Will gladly come back!!



Rick Uwens

While the kitchen is highly competent and the staff very obviously know what they are doing, I must unfortunately subtract a single star: On occasion and very depending on the batch and presumably the exact person taking care of said batch, the fries tend to be slightly soggy.
It does not take away from the overwhelmingly beautiful interieur or the otherwise exceptional menu though. Recommended! :-)



Anne A. Littics

Absolute recommendation! Kids and I love it here!



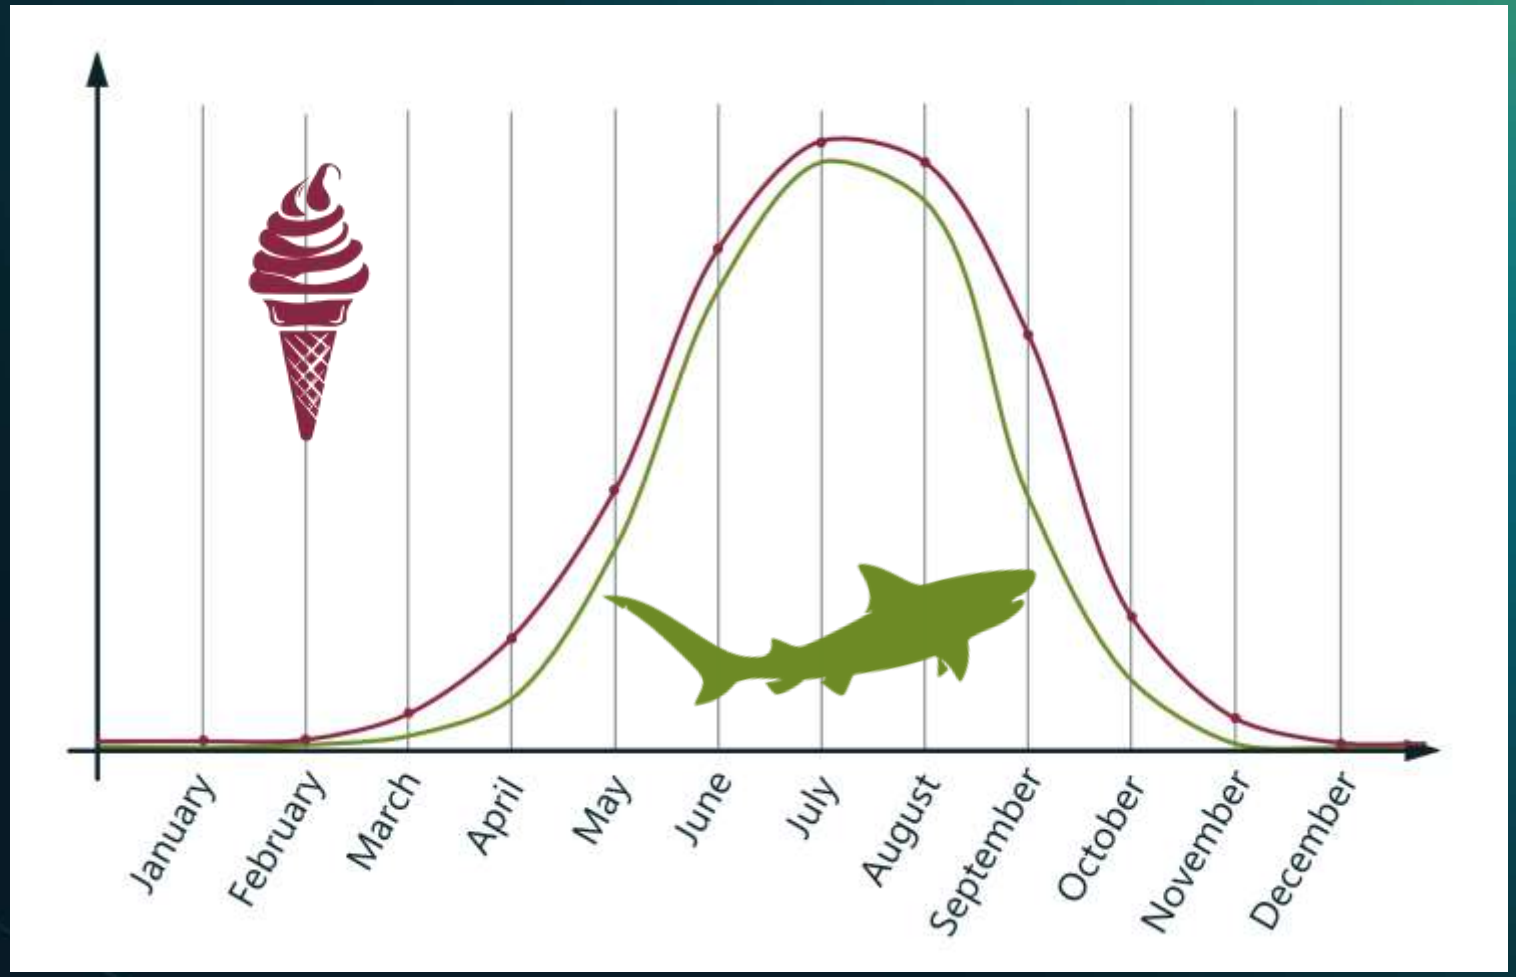
Mike Rusoff

Competent staff and cooking. Special recommendation for the potato wedgies, and ALL the housemade sauces.
Easily the best burger place in town, plus the open kitchen is always a highlight

In a **day-to-day**, data-driven environment, data literacy empowers us to explore and analyze data **by depending** on our critical thinking skills.



In a **day-to-day**, data-driven environment, data literacy empowers us to explore and analyze data **by depending** on our critical thinking skills.





In a **day-to-day**, data-driven environment, data literacy empowers us to explore and analyze data **by depending** on our critical thinking skills.



Source: If Google Was A Guy

Why does Data Literacy matter?

- Spot **biases** or manipulation
  - Not only in the report, but **in yourself as well**
- Judge the **credibility** of sources
- Be able to act and decide by **interpreting** everyday data **correctly**

All rooted in the ability to **think critically** and **research meaningfully**.

So how do you measure it?

Is everyone in your company data literate?

Are people in your personal environment data literate?

How and where do you start measuring it?

- People must be verbally, numerically, and graphically literate
- How to measure abstract values like critical thinking?

So how do you measure it?

1. Self-assessment

2. Normed quizzes

3. Observed behavior

1. Self-assessment

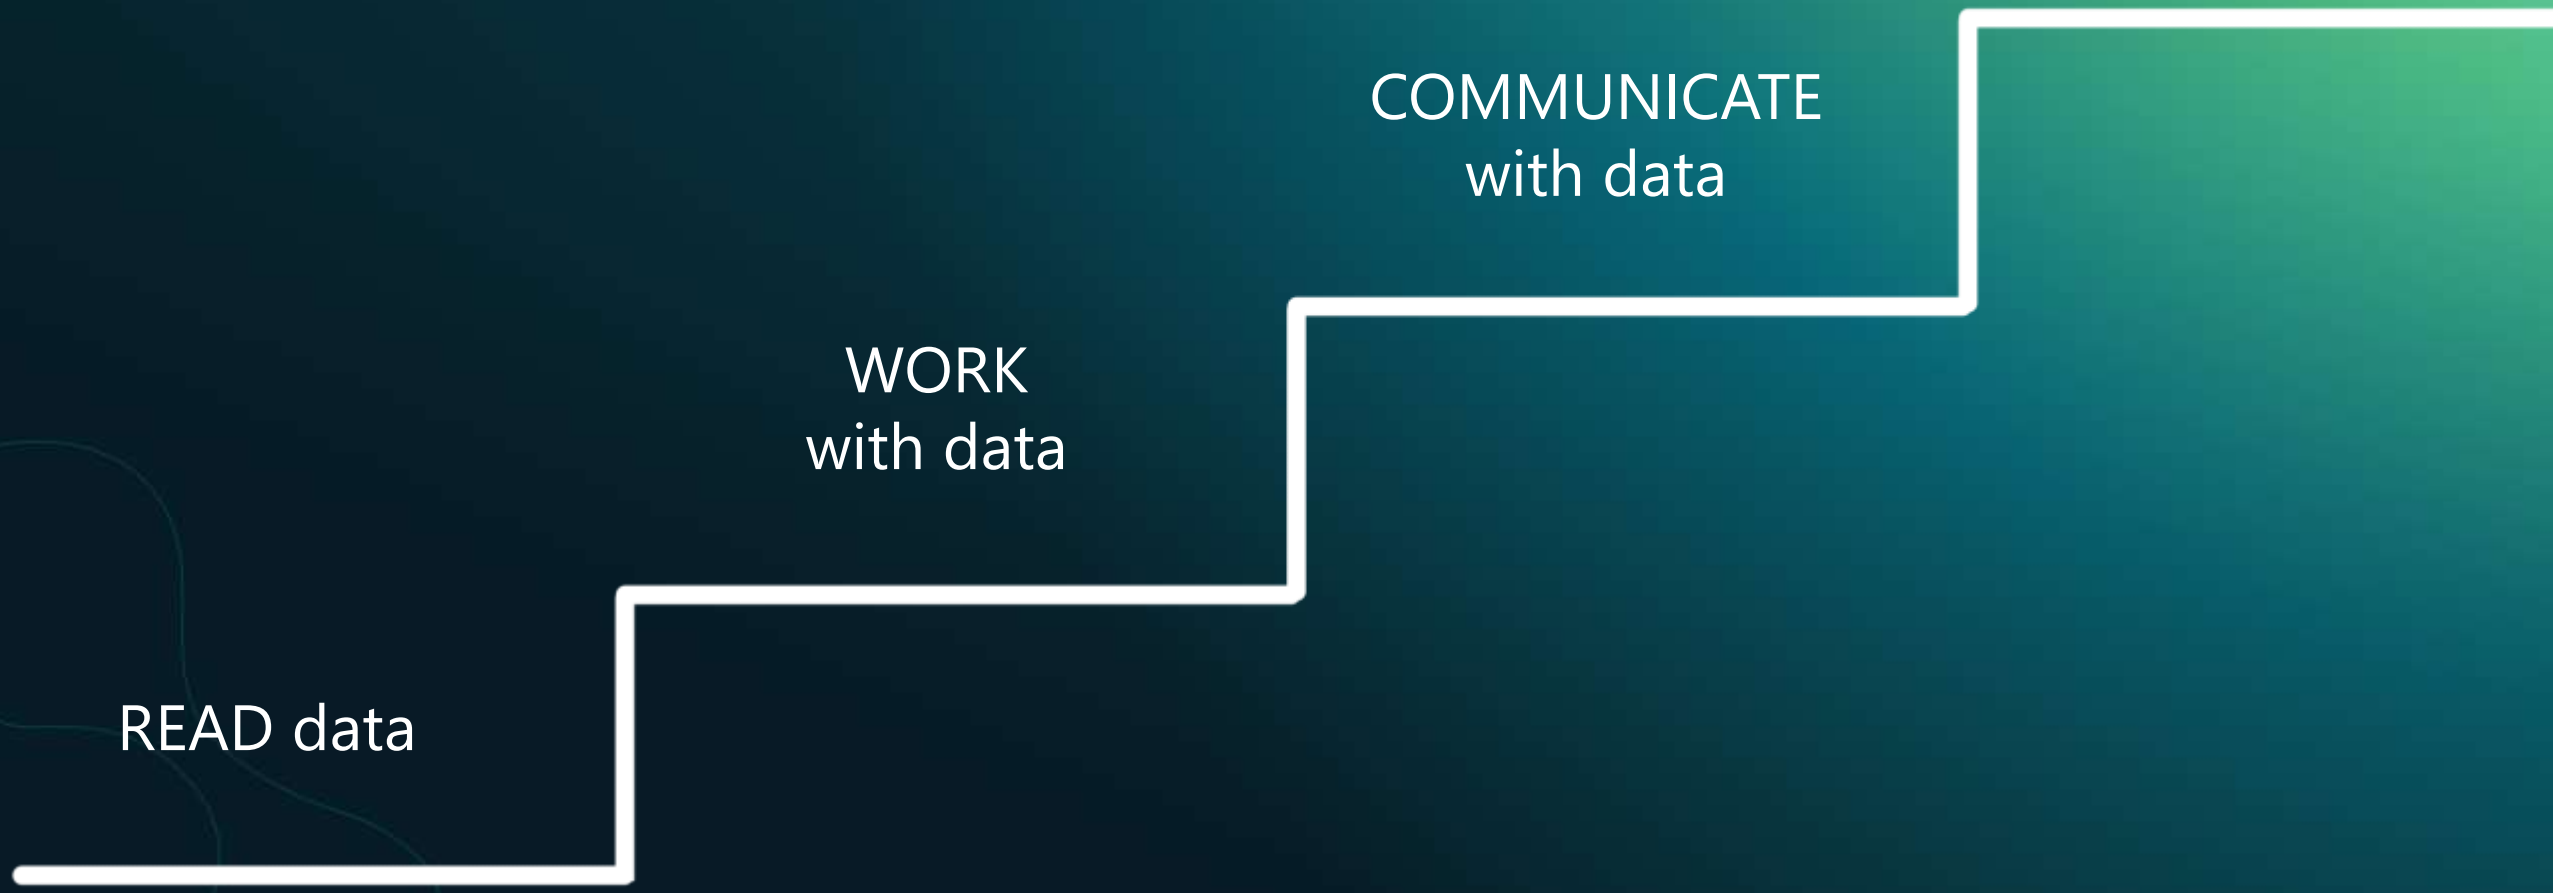
- This is always a highly **subjective** matter!
    - Self-esteem, personality, environment play an influential role

- Most valuable in **context**:
    - Which role does the employee occupy?

- Measurement **over time**:
    - Which impact has training had?
    - Did it improve confidence?

1. Self-assessment

What to evaluate?



READ data	WORK with data	COMMUNICATE with data	DECIDE with data
- Discover - Evaluate - Interpret	- Collect - Analyze - Transform - Curate - Manage	- Interpret - Visualize - Present	- Innovate - Conclude - Drive decisions

How to evaluate?

Competency Rating Scale:

Level 1	I can finish **simple tasks** with clear **instruction/ guidance**
Level 2	I can finish tasks **on my own**
Level 3	I can solve **problems on my own**
Level 4	I can **assist** others with simple tasks and problems
Level 5	I can **teach** others to solve complex tasks and problems

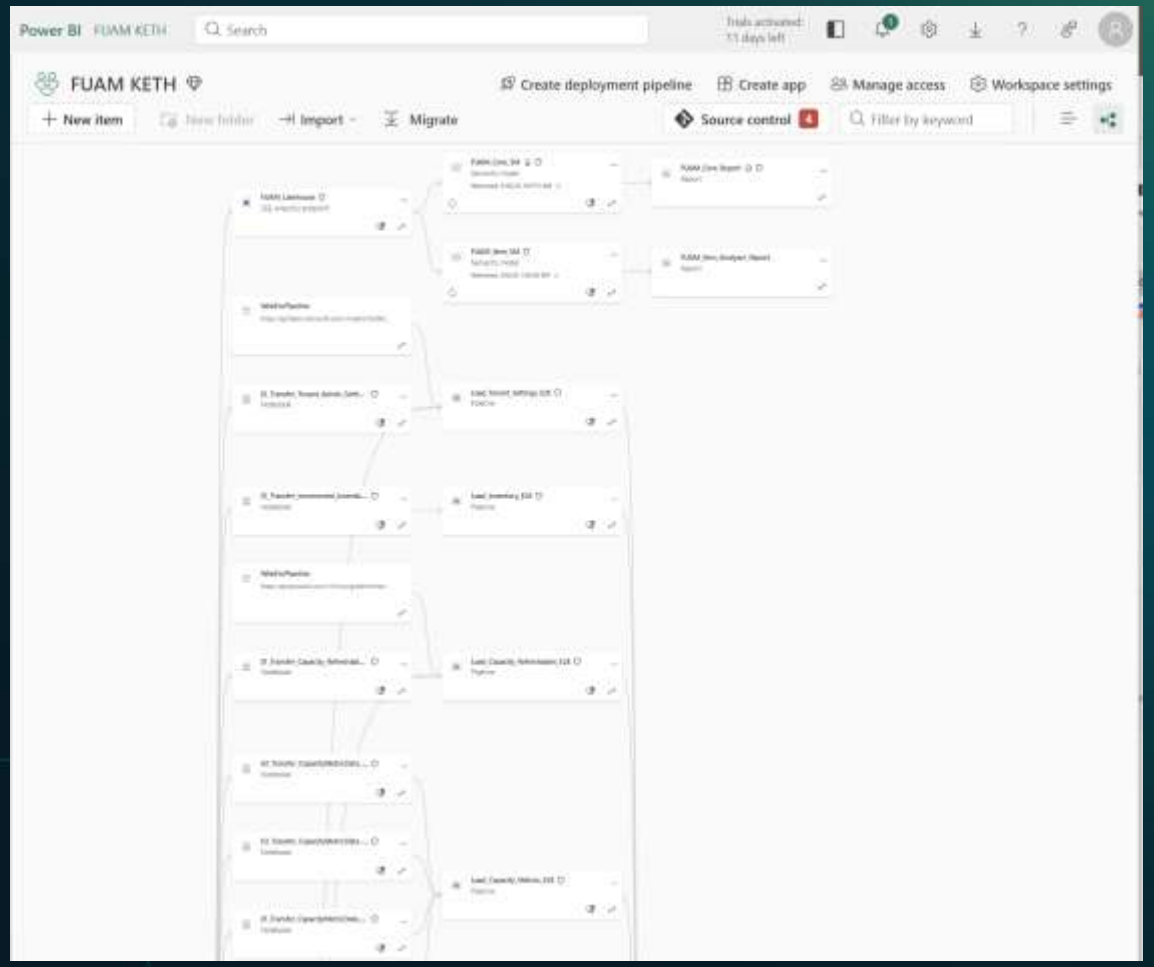
Discover data

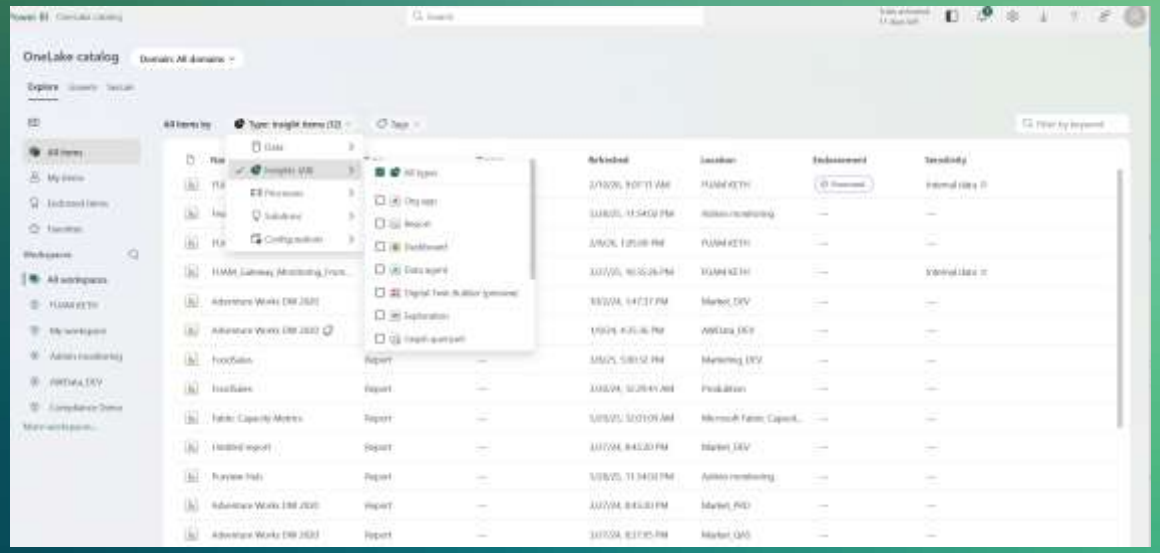
Level 1	Level 2	Level 3	Level 4	Level 5
I can **open** a specified dataset/ report in a location **provided to me**	I can **browse** and access the data I need from a range of sources	I can **search** various sources to find data related to my needs	I can **help others** search a range of sources related to their needs	I can **teach others** how to find and access the data they need

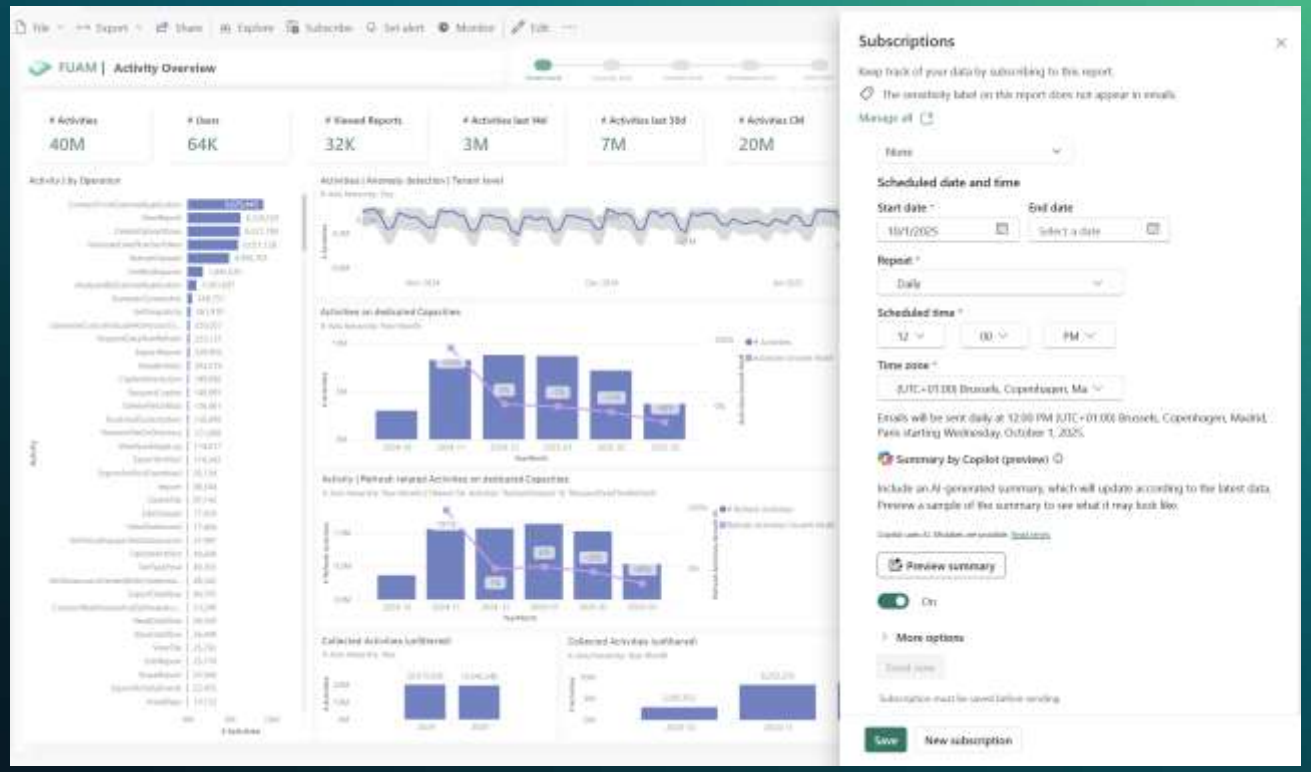


Analyze data

Level 1	Level 2	Level 3	Level 4	Level 5
I can **use** simple tools **provided to me** to undertake analysis	I can **use** the most appropriate tool for my analysis	I can **research/ create** the most appropriate tool for my analysis	I can **help others** select and use the most appropriate tools	I can **teach others** which tools are most appropriate and why


Power BI
FUAM KETH
Search
FUAM KETH
Create deployment pipeline
Create app
Manage access
Workspace settings
New item
Import
Migrate
Source control
Filter by keyword


OneLake catalog


FUAM | Activity Overview
Activities
40M
Users
64K
Viewed Reports
32K
Activities last Wk
3M
Activities last 30d
7M
Activities Old
20M
Subscriptions
Keep track of your data by subscribing to this report.
Scheduled date and time
Start date
End date
Repeat
Daily
Scheduled time
Time zone
Summary by Copilot (preview)
Preview summary
More options
Save
New subscription

Visualize data

Level 1	Level 2	Level 3	Level 4	Level 5
I can **use templates** to create tables and charts presenting data provided to me	I can **create** meaningful tables and charts to present a dataset	I can **visually present** data from a range of sources meaningfully	I can **help others** visually present simple data sources	I can **teach others** how to visually present data meaningfully



Question data

Level 1	Level 2	Level 3	Level 4	Level 5
I can **answer** basic questions given with data provided to me	I can **find answers** regarding a stated question in a range of data sources	I **know what to ask** and where to **find answers** regarding a broader issue	I can **assist others** in pointing out what to ask and where to find answers	I can **teach others** how to ask the right questions and where to find answers

2. Quizzes/ Tests

- Can raise **awareness**
    - About unknown deficiencies or about topics/ issues in general

- Can already be tailored to emphasis **specific topics or accommodate roles**

- Measurement **over time**:
    - Did training measures succeed?

2. Quizzes/ Tests

Question

A chart looks entirely unreasonable and confusing
What do you do?

- A: Analyze the underlying data carefully
- B: Blame your boss
- C: Blame society
- D: Post it on Slack and hope someone else fixes it

Question

A dataset updates nightly.
What does this imply?

- A: The data is unreliable
- B: It updates manually
- C: Today's decisions use yesterday's data
- D: It is nearly realtime

Question

A dashboard shows unexpected results.
What is the most appropriate first step?

- A: Use a Pie chart instead
- B: Check the data source and transformation steps
- C: Refresh the page
- D: Ignore the dashboard

Question

What does SQL stand for?

- A: Standardized Query Language
- B: Structured Query Language
- C: Slowly Questioning Leadership
- D: System Query Language

2. Quizzes/ Tests

Question What does the abbreviation 'SHIR' stand for?

- A: Save Hours with Incremental Refresh
- B: SHit I Ran out of coffee
- C: Some Humans Ignore Reality
- D: Stuff Happens in Reports

Question Which SQL statement retrieves all columns from a table named Sales?

- A: GET * FROM Sales
- B: SHOW ALL * FROM Sales
- C: SELECT * FROM Sales
- D: EXTRACT * FROM Sales

Question Why is scripted analysis (SQL, Python) preferred over manual steps?

- A: It looks more professional
- B: It is easier to automate and reproduce
- C: It requires less thinking
- D: It avoids documentation

Question You join two tables and rowcount suddenly doubles. What is the most likely cause?

- A: Data corruption
- B: Many-to-many join
- C: Incorrect data types
- D: Missing values

2. Quizzes/ Tests

- Relevancy to the role/ business
- Adjusting difficulty
- Best usecase: check if implemented training was successful

3. Observed behavior

= Monitor usage trends

- In a **self-service setting** like Fabric/ Power BI → measuring Data Literacy through **interaction**
- FUAM to **track adoption**
- Monitor **report performances**

- **! Important** to have an established **governance foundation**, allowing everyone easily understandable and secure access
    - Define and publish **certified datasets**
    - Provide **templated reports**
    - Implement **RLS** beforehand

3. Observed behavior

= Monitor usage trends

- Measure data-literacy maturity in a self-service setting needs **Telemetry that Fabric / Power BI already captures**
- The program should be flexible and **measurement should apply to work-council and other regulatory standards**

3. Observed behavior

= Monitor usage trends

Ability	Typical behaviours	Examples of Power BI signals
Read	Finds certified data, uses filters, interprets visuals correctly	Report Views, Distinct Pages Viewed
Work	Adds slices, drills, creates personal bookmarks, builds reports, built semantic models, write DAX, publishes to workspace	ViewReport + QueryDataset events, CreateReport, CreateDataset events
Communicate & Share	Shares links, schedules subscriptions, comments	ShareDashboard, CreateSubscription
Governance habits	Uses certified datasets, applies sensitivity labels	Dataset lineage + label activity

3. Observed behavior

= Capture quantitative evidence with standard Fabric telemetry

Source	What it gives you	How to enable / query
Usage Metrics report (per-item & workspace)	90-day counts of views, shares, retention curves, device split	Available to every workspace admin → **View usage metrics report** **More advanced with FUAM**
Tenant Activity Log (Audit)	Raw events such as ViewReport, CreateReport, UpdateDataset, ShareDashboard with actor, timestamp, workspace	Pull via **Admin – Get Activity Events** REST API or PowerShell Get-PowerBIActivityEvent; up to 30 days per More advanced on the FUAM Lakehouse Data
Adoption-tracking guidance	Possible KPIs, Measures	Microsoft Adoption Playbook
Fabric adoption roadmap	Maturity definitions & benchmark questions	Helps translate telemetry into laddered maturity levels

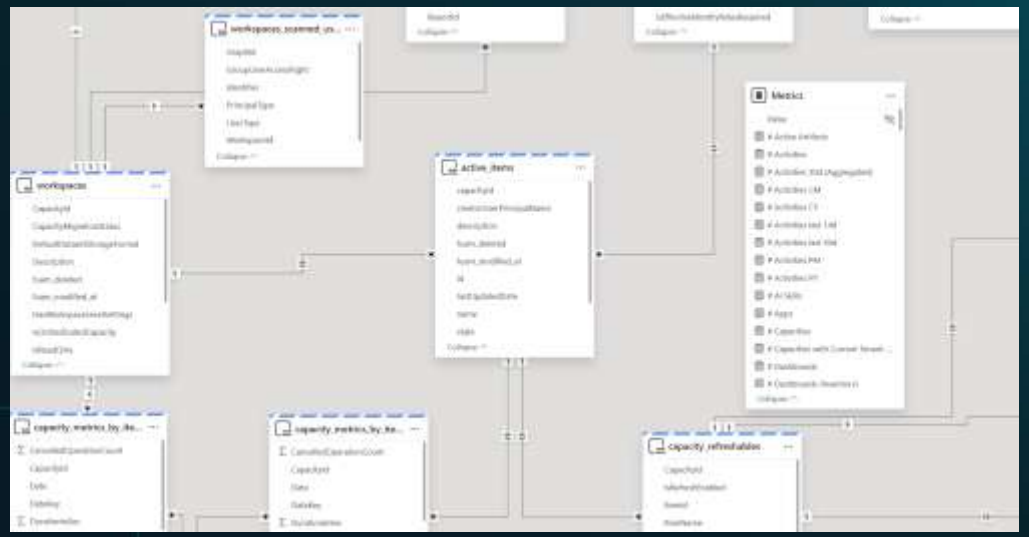
3. Observed behavior

= Capture quantitative evidence with standard Fabric telemetry

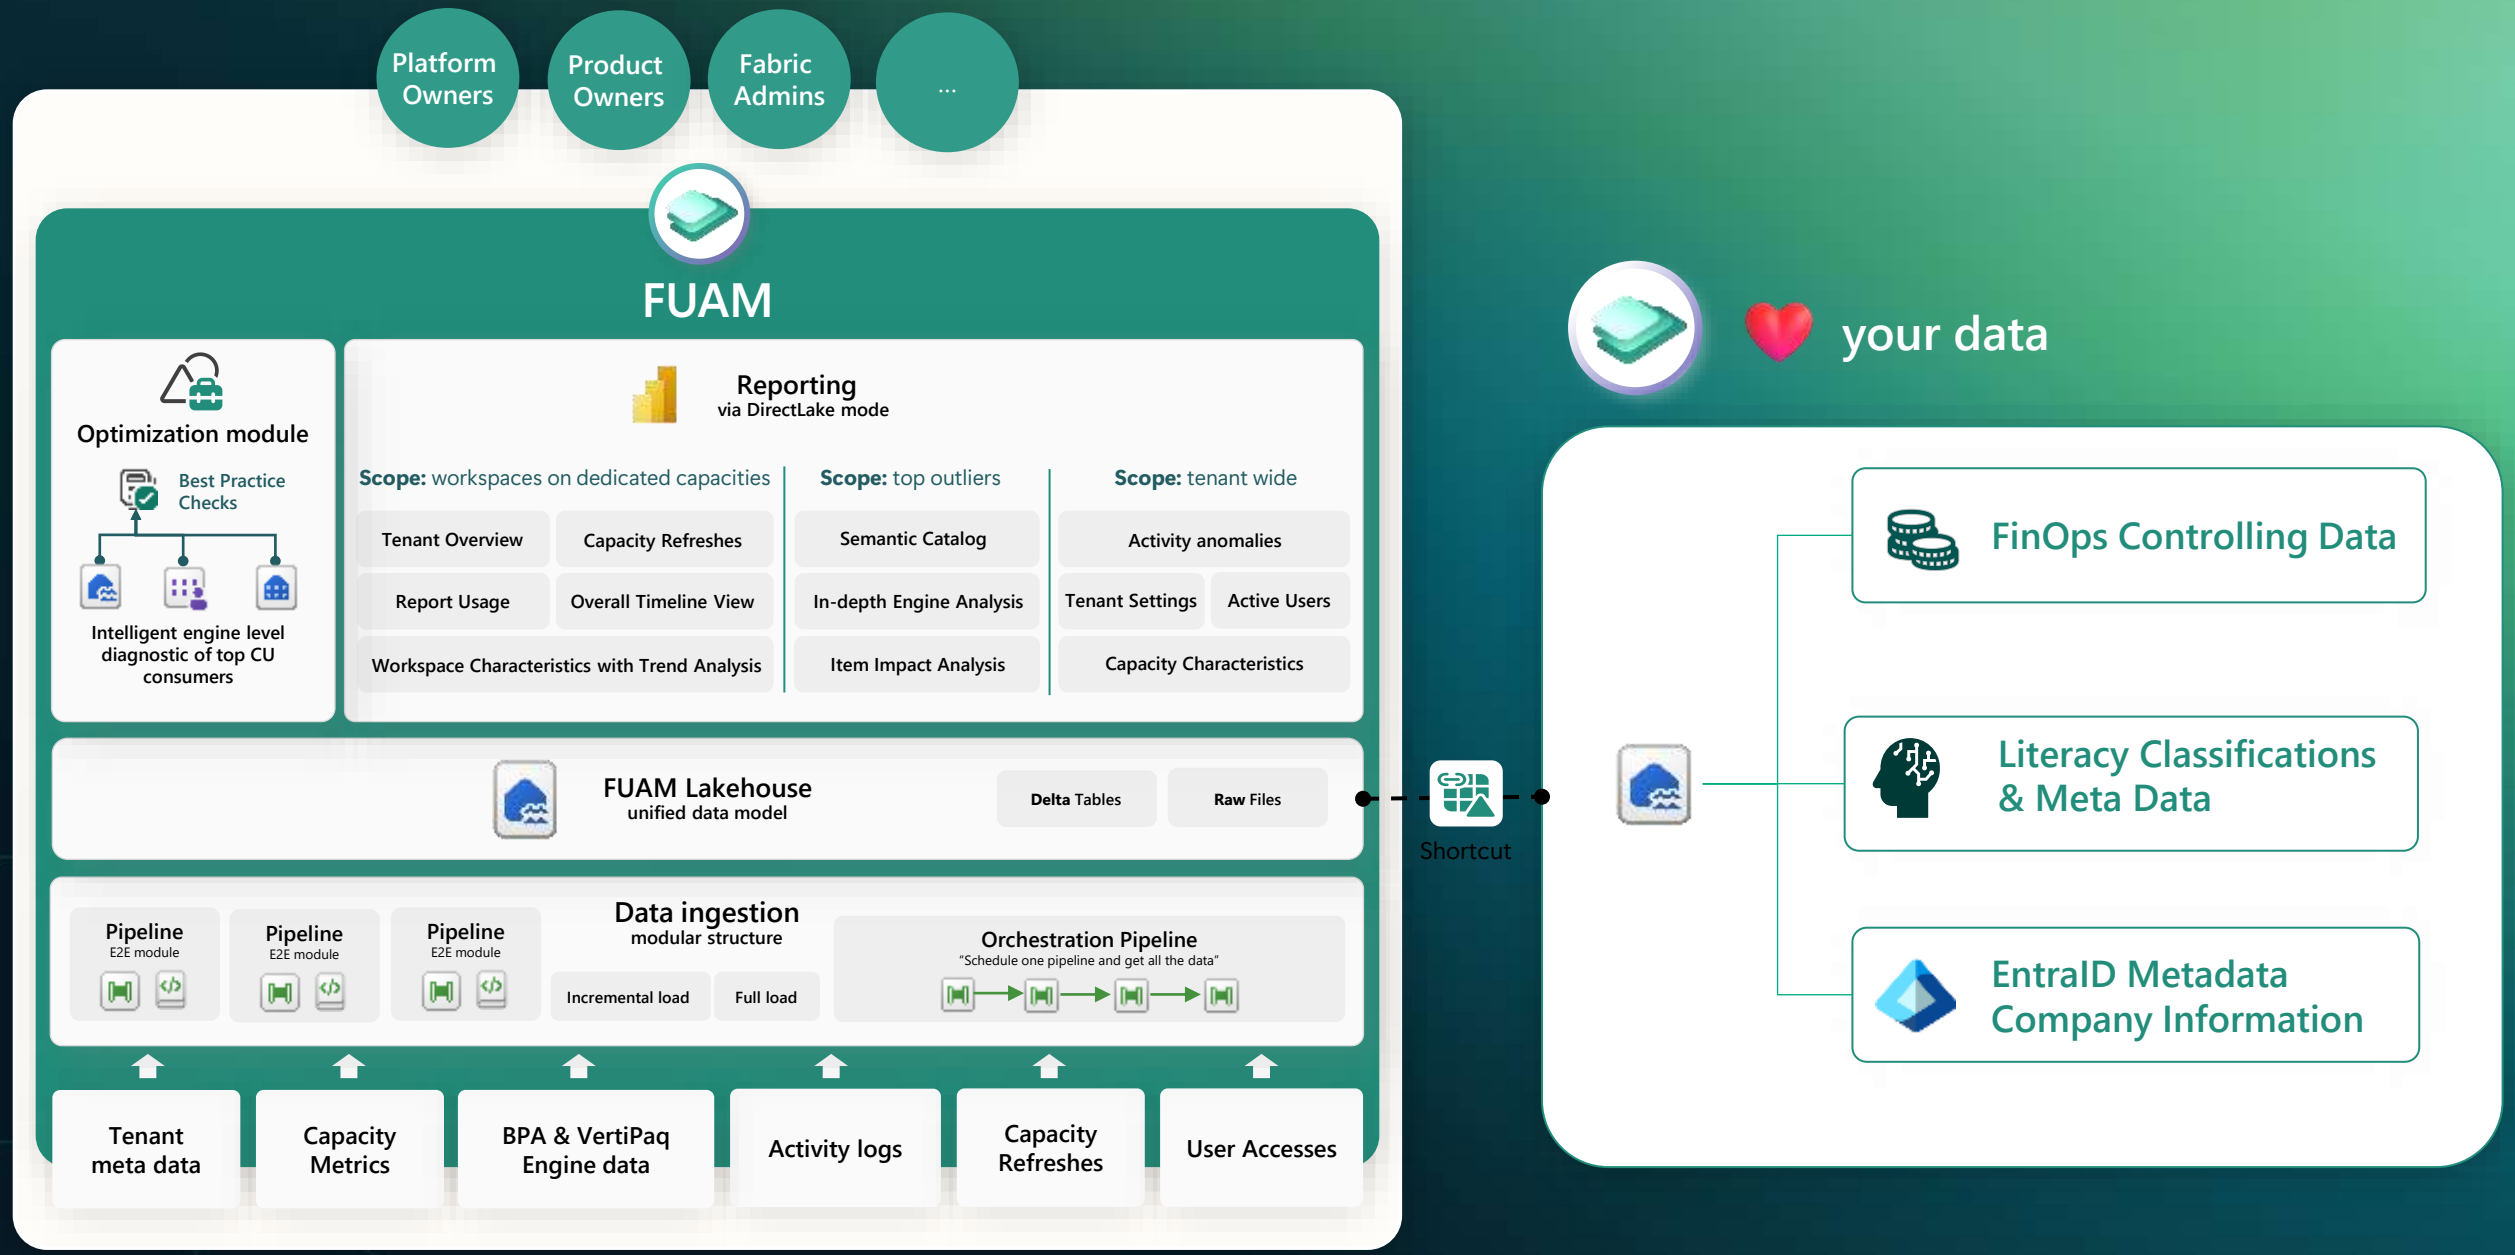
Component	Typical Power BI signals	Why it matters
Active Creators (30 %)	• Distinct users who triggered CreateReport, CreateDataset, UpdateDataset, PublishApp in the Activity Log during the period • Usually expressed as a **percentage of licensed users** in the same business unit to stay size-neutral	Captures the breadth of self-service authorship—people moving from consumption to creation
Depth of Interacti on (20 %)	• Median drill-through actions per session (DrillThroughReport events) • Avg. filters/slicers applied • % of sessions with > *1* page viewed	Indicates analytical “stickiness”: users don’t just glance at a KPI—they explore
Sharing Index (15 %)	Weighted sum of: • ShareReport and ShareDashboard events (x 1) • New subscriptions (CreateSubscription) (x 1.5) • Mentions/comments via Teams integration (x 2)	Collaboration is a strong predictor of data-driven decision-making culture
Governance Com pliance (15 %)	• % of reports built on **certified** or **promoted** datasets • % of artefacts with a **sensitivity label** • % workspaces in the **Production** capacity tier	Ensures self-service growth doesn’t erode trust, security or lineage

3. Observed behavior

= Capture quantitative evidence with standard Fabric telemetry






Platform Owners
Product Owners
Fabric Admins
...
FUAM
Optimization module
Best Practice Checks
Intelligent engine level diagnostic of top CU consumers
Reporting
via DirectLake mode
Scope: workspaces on dedicated capacities
Tenant Overview
Capacity Refreshes
Report Usage
Overall Timeline View
Workspace Characteristics with Trend Analysis
Scope: top outliers
Semantic Catalog
In-depth Engine Analysis
Item Impact Analysis
Scope: tenant wide
Activity anomalies
Tenant Settings
Active Users
Capacity Characteristics
FUAM Lakehouse
unified data model
Delta Tables
Raw Files
Shortcut
Pipeline
E2E module
Pipeline
E2E module
Pipeline
E2E module
Data ingestion
modular structure
Incremental load
Full load
Orchestration Pipeline
"Schedule one pipeline and get all the data"
Tenant meta data
Capacity Metrics
BPA & VertiPaq Engine data
Activity logs
Capacity Refreshes
User Accesses
your data
FinOps Controlling Data
Literacy Classifications & Meta Data
EntraID Metadata Company Information

And how do you improve it?

- Define clear goals and metrics
- Raise awareness
    - Communicate the values of data literacy
    - “Data days”, webinars, newsletters
- Training and enablement
    - Tiered learning paths based on user roles (workshops, internal communities, train-the-trainer models, …)

Who should take care of it?

- Gap of expectations
- Responsibility to learn vs. responsibility to teach?
- Everybody expects the other to do it



Cost and Investment

"We already invested in BI tools and platforms."

"Training everyone is not cost-effective."

"Learning data skills on my own costs **my** time and energy."

Time and Capacity

"We can't pull teams out of their daily work."

"Learning should happen on the job."

"There's no protected time to practice."

"Learning data skills feels like extra unpaid work."

Responsibility

"Data literacy is a universally required skill nowadays."

"People don't engage with the training we offer."

"If it's important, why isn't it formalized?"

"We are expected to be data-driven but not taught how."

Motivation and Confidence

"You can't teach motivation."

"Not everyone wants to work with data."

"Numbers look intimidating and I will look incredibly stupid."

Forrester Consulting – Building Data Literacy – February 2022

82%

Of **leaders** expect basic data literacy from all employees

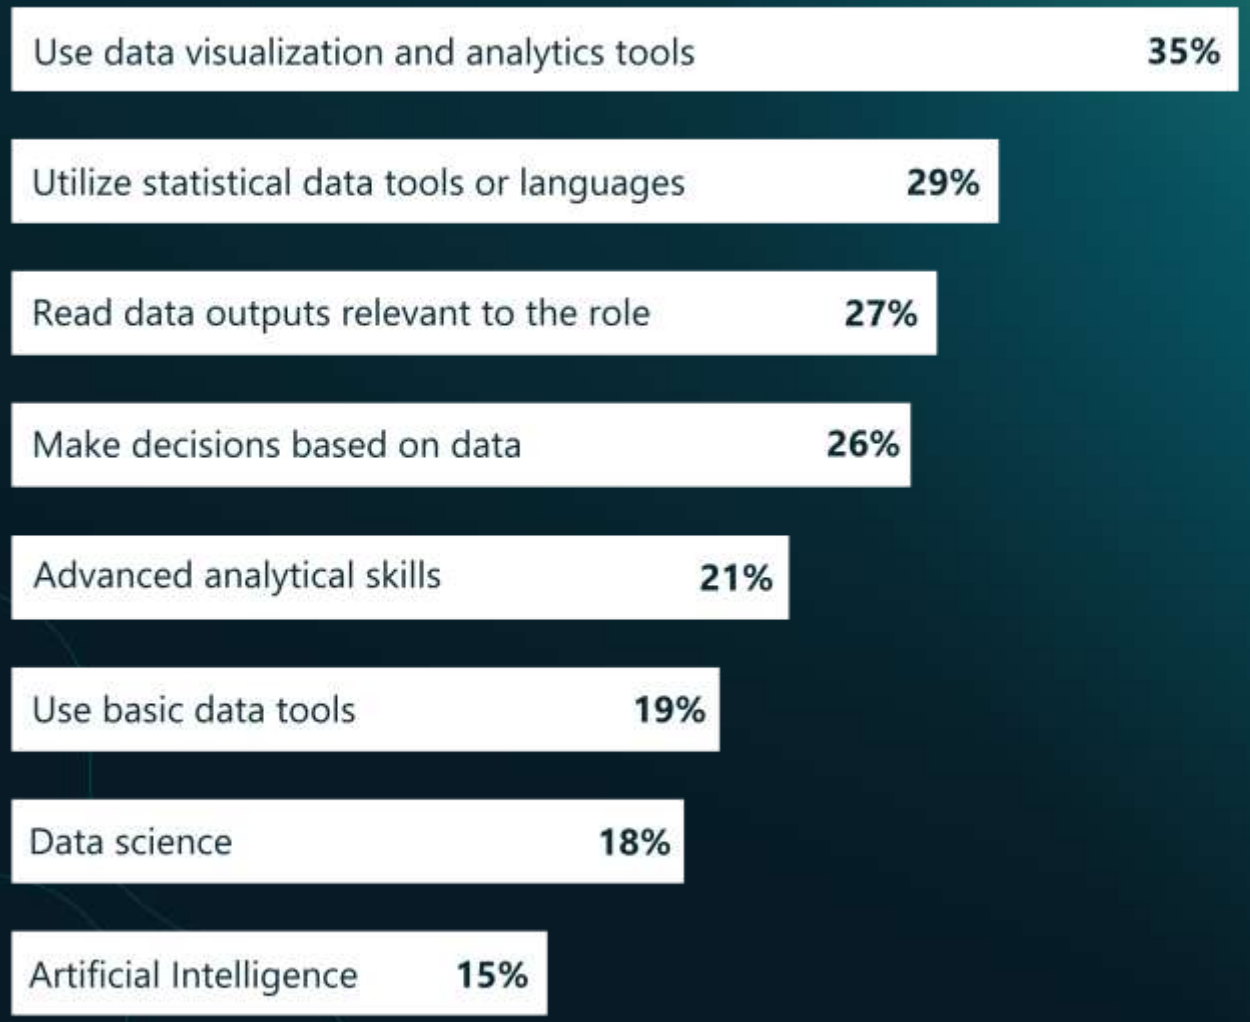
47%

Of **employees** say that they have been offered data training

Source

47%

Of employees say that they have been offered data training



79%

Of **leaders** say their department is successfully equipping their employees with the data skills they are expected to have

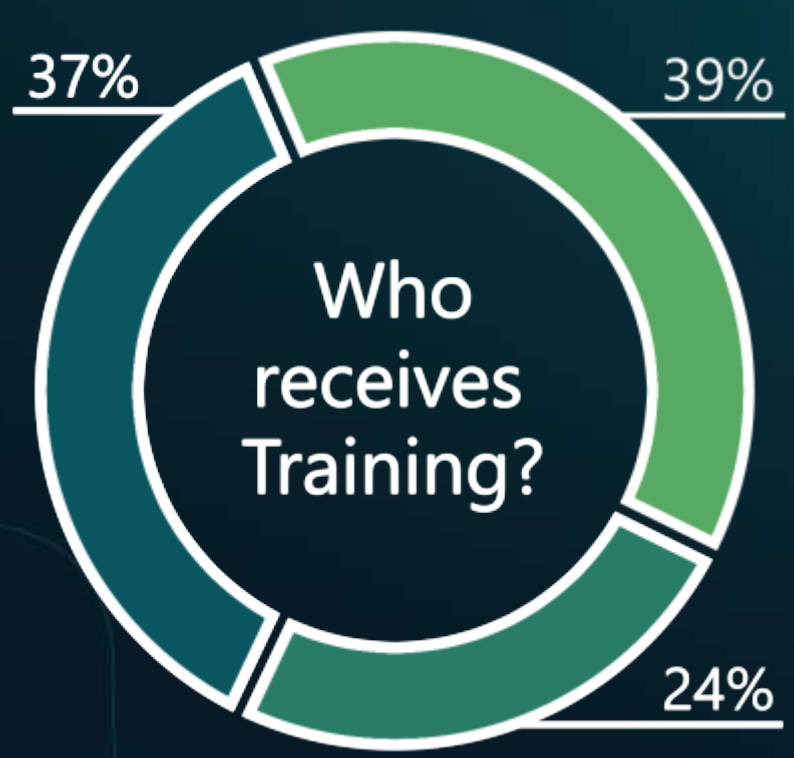
40%

Of **employees** say their organization has successfully equipped them with the data skills they are expected to have

Source

47%

Of **employees** say that they have been offered data training





Source

% of employees who think Data Training is '*important*' or '*very important*'



Analytics
85%



Finance
80%



Marketing
87%



HR
86%



Research
83%



IT
90%

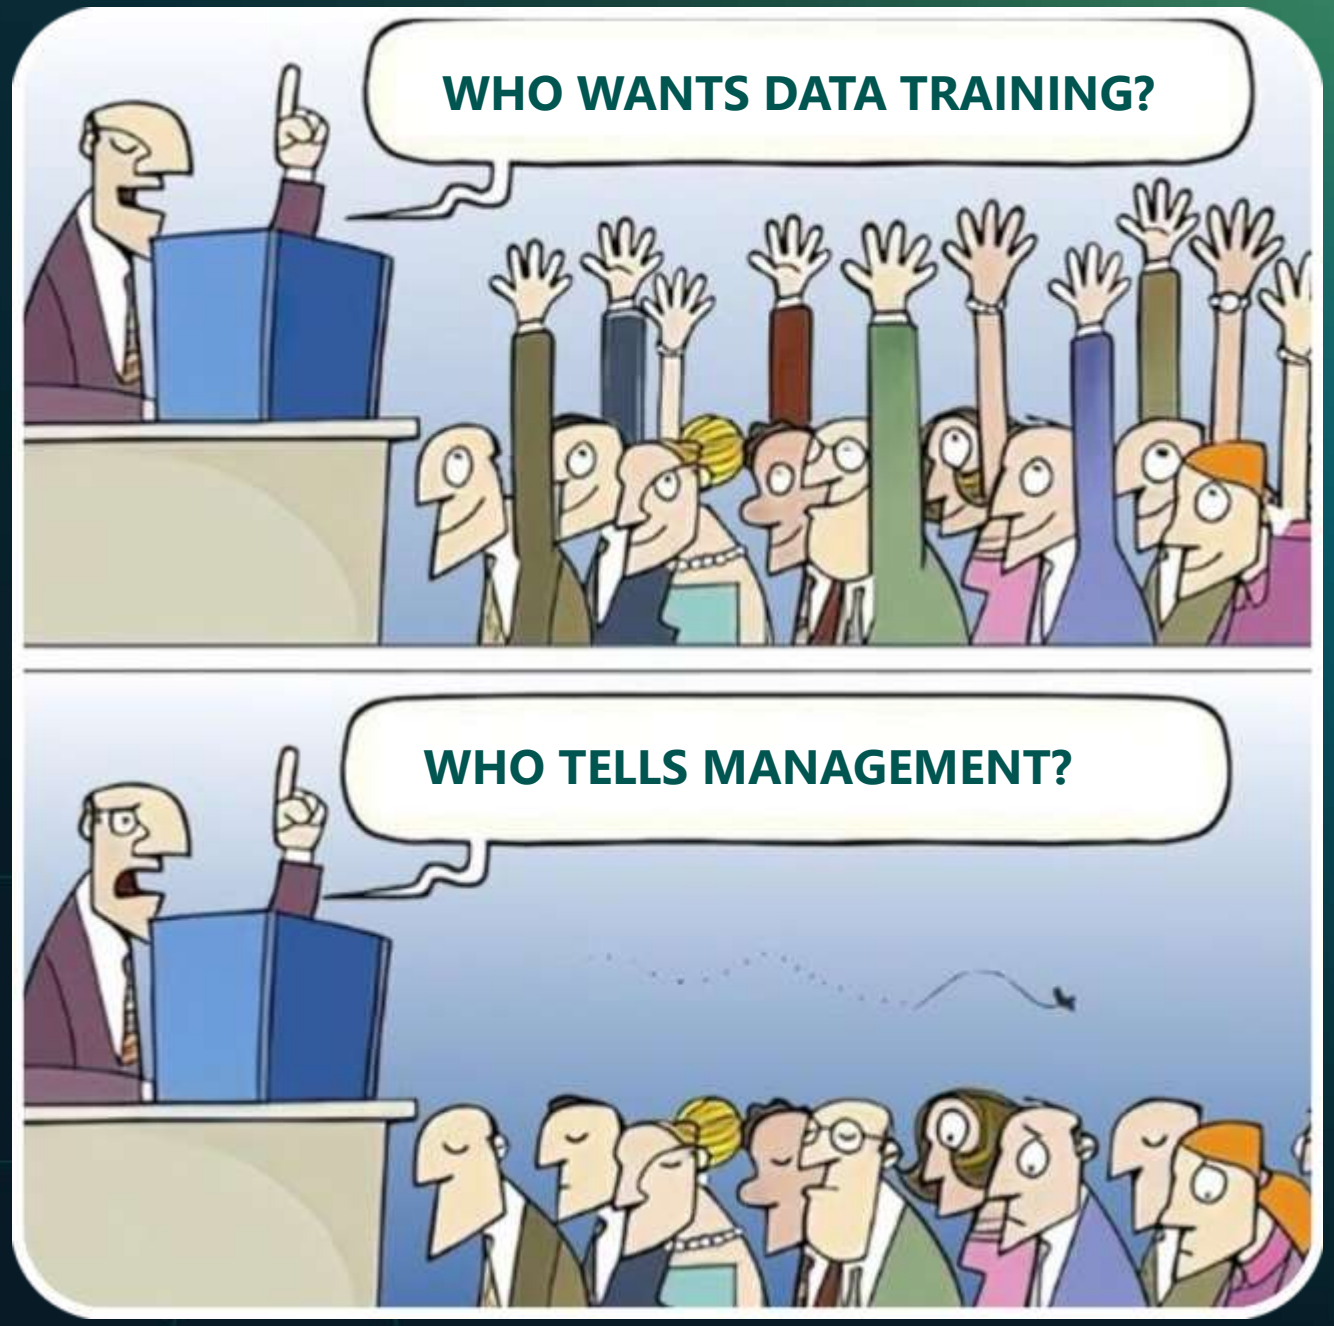


Operations
86%



Engineering
81%

Source



26%

Of employees say that they talked to management about wanting more data training.

16%

Of employees say that they complained about the amount of data training offered.

Source



Lack of skilled staff to lead training



Lack of budget



Lack of knowledge on how to improve data skills



Lack of support from the wider organization



Lack of a data culture



Resistance from the employees



Inadequate training materials



Lack of skilled staff to lead training



Lack of budget



Lack of knowledge on how to improve data skills



Lack of support from the wider organization



Lack of a data culture



Resistance from the employees



Inadequate training materials

How to solve this?

Governance

- Make things **easy to access** and understand
- Ensure **data quality**
- Utilize **embedded tools** like Power BI descriptions, tooltips, formatting
- Maintain easy to find and read data **glossaries and documentations**

→ Guided **Self-Service**

→ Establish **confidence**

How to solve this?

Training and Enablement

- Leadership **commitment**
    - **Communicate** the value of Data Literacy
    - Data-days, webinars, newsletters

- Train-the-Trainer, Lunch-and-Learn models

- Tiered learning paths, based on role, create **relevancy**

→ Keep **Motivation** high

How to solve this?

Tool Familiarity

- Smaller onboarding sessions introducing Power BI/ Fabric
- Create a defined system for support, practice communities

Gamification

- Not a competition!
- But value top contribution
- System for badges, certifications for completing certain milestones

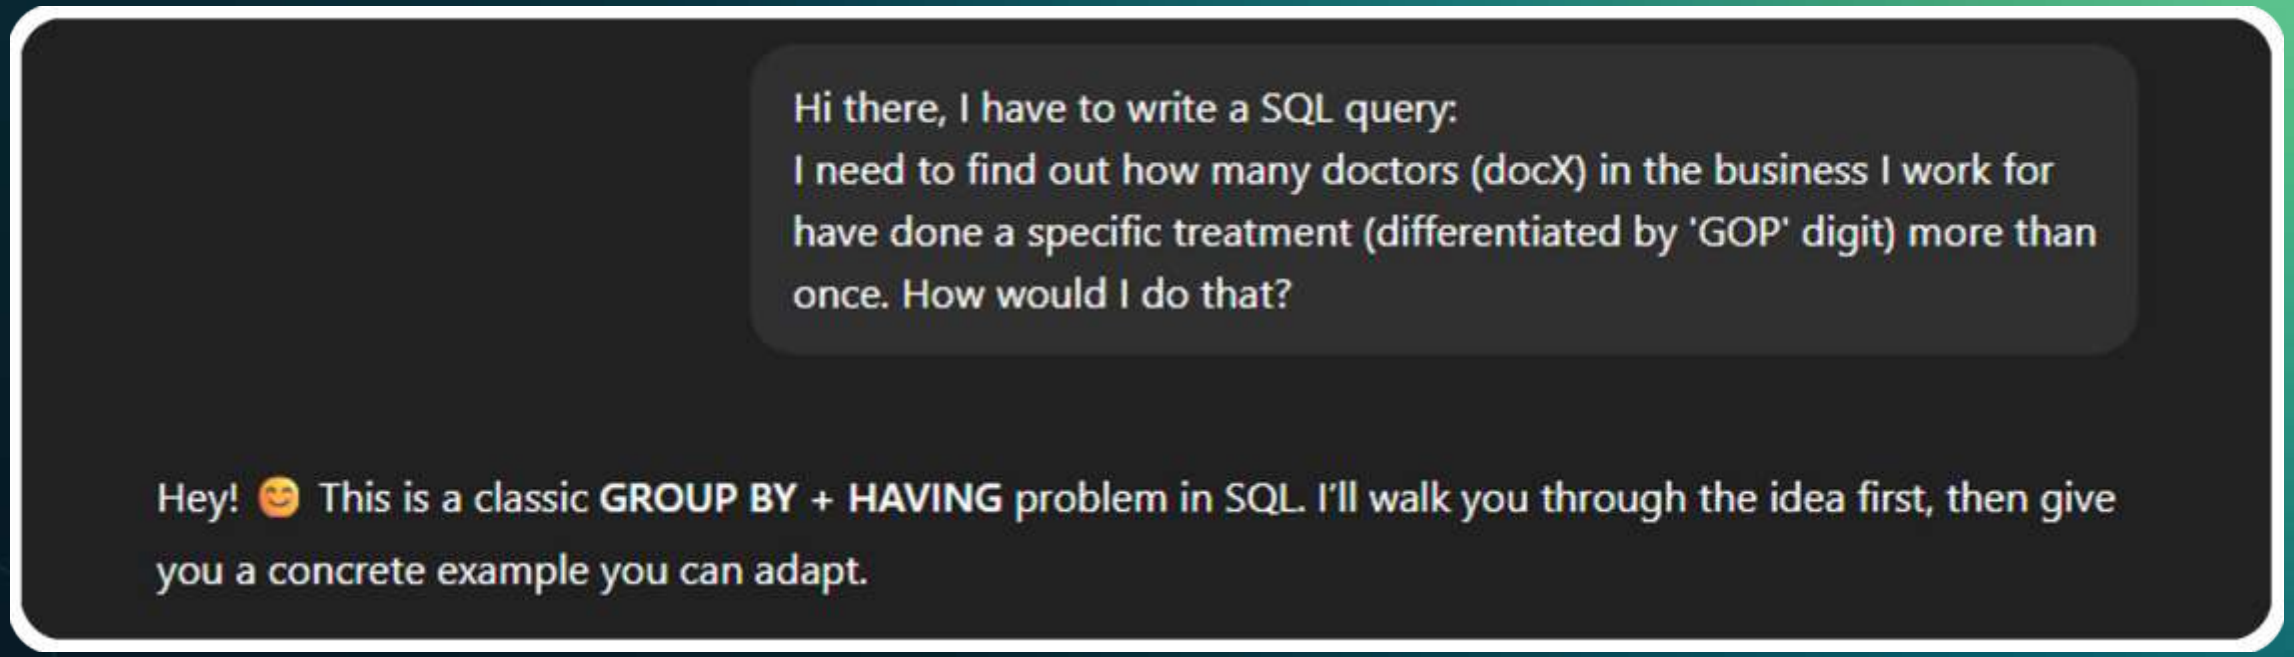
AI and Data Literacy

- Understand LLMs
    - Where they receive information from
    - How they are coded

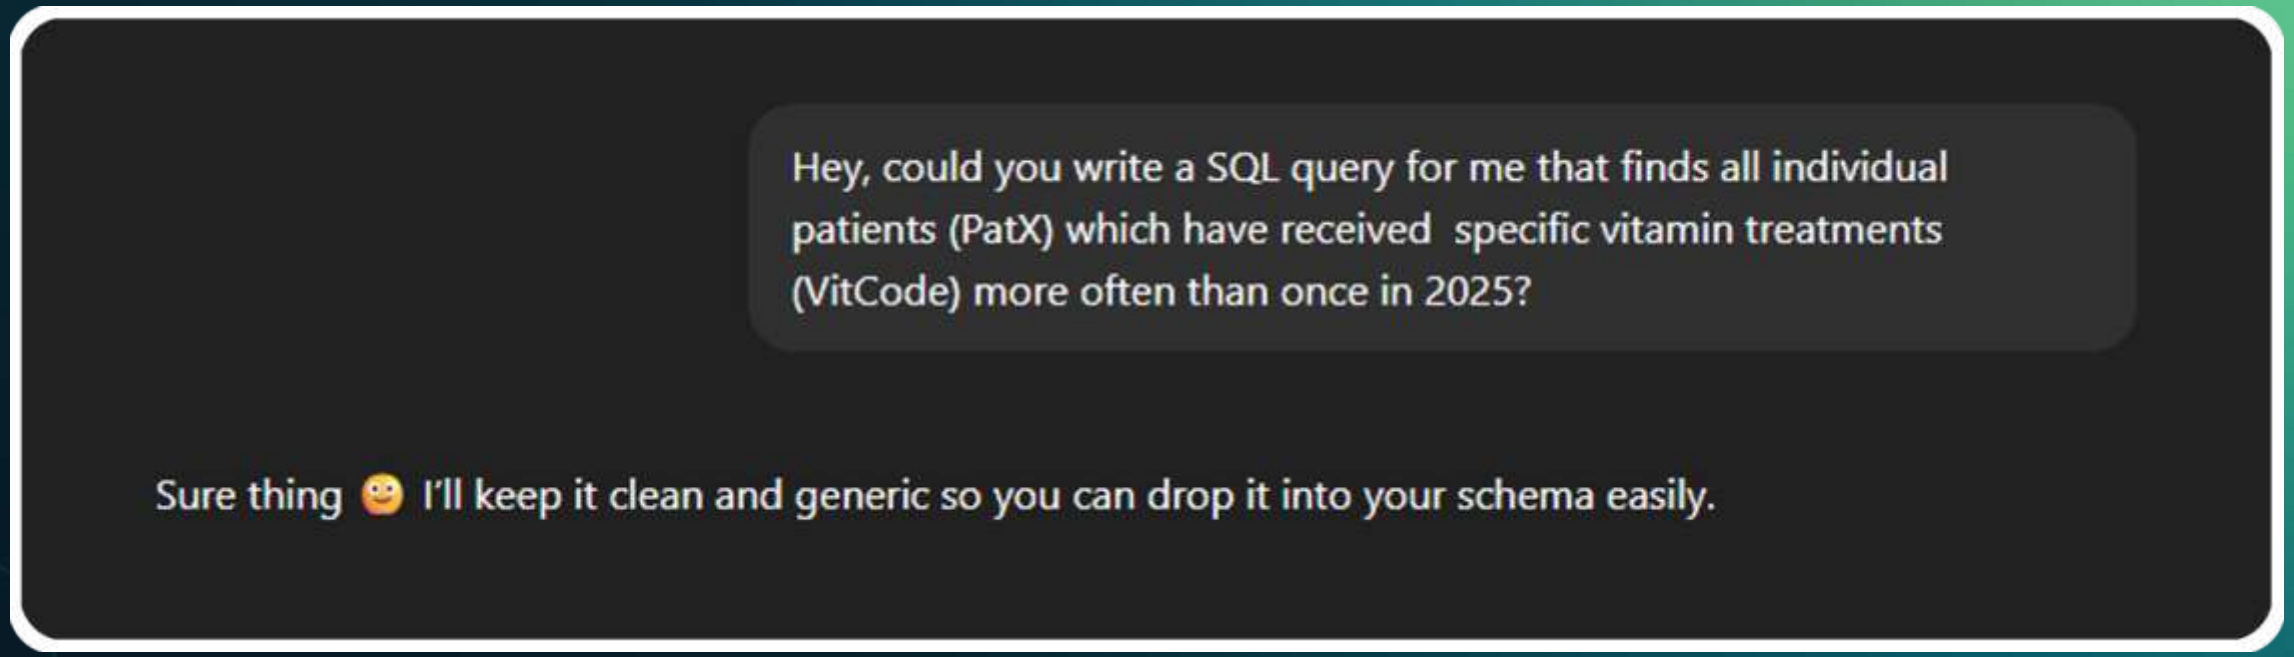
- Write effective prompts, skill to phrase good prompts
- Correctly interpret a reply

- Potential inaccuracies and hallucinations

But that’s not the point.



But that’s not the point.



But that’s not the point.

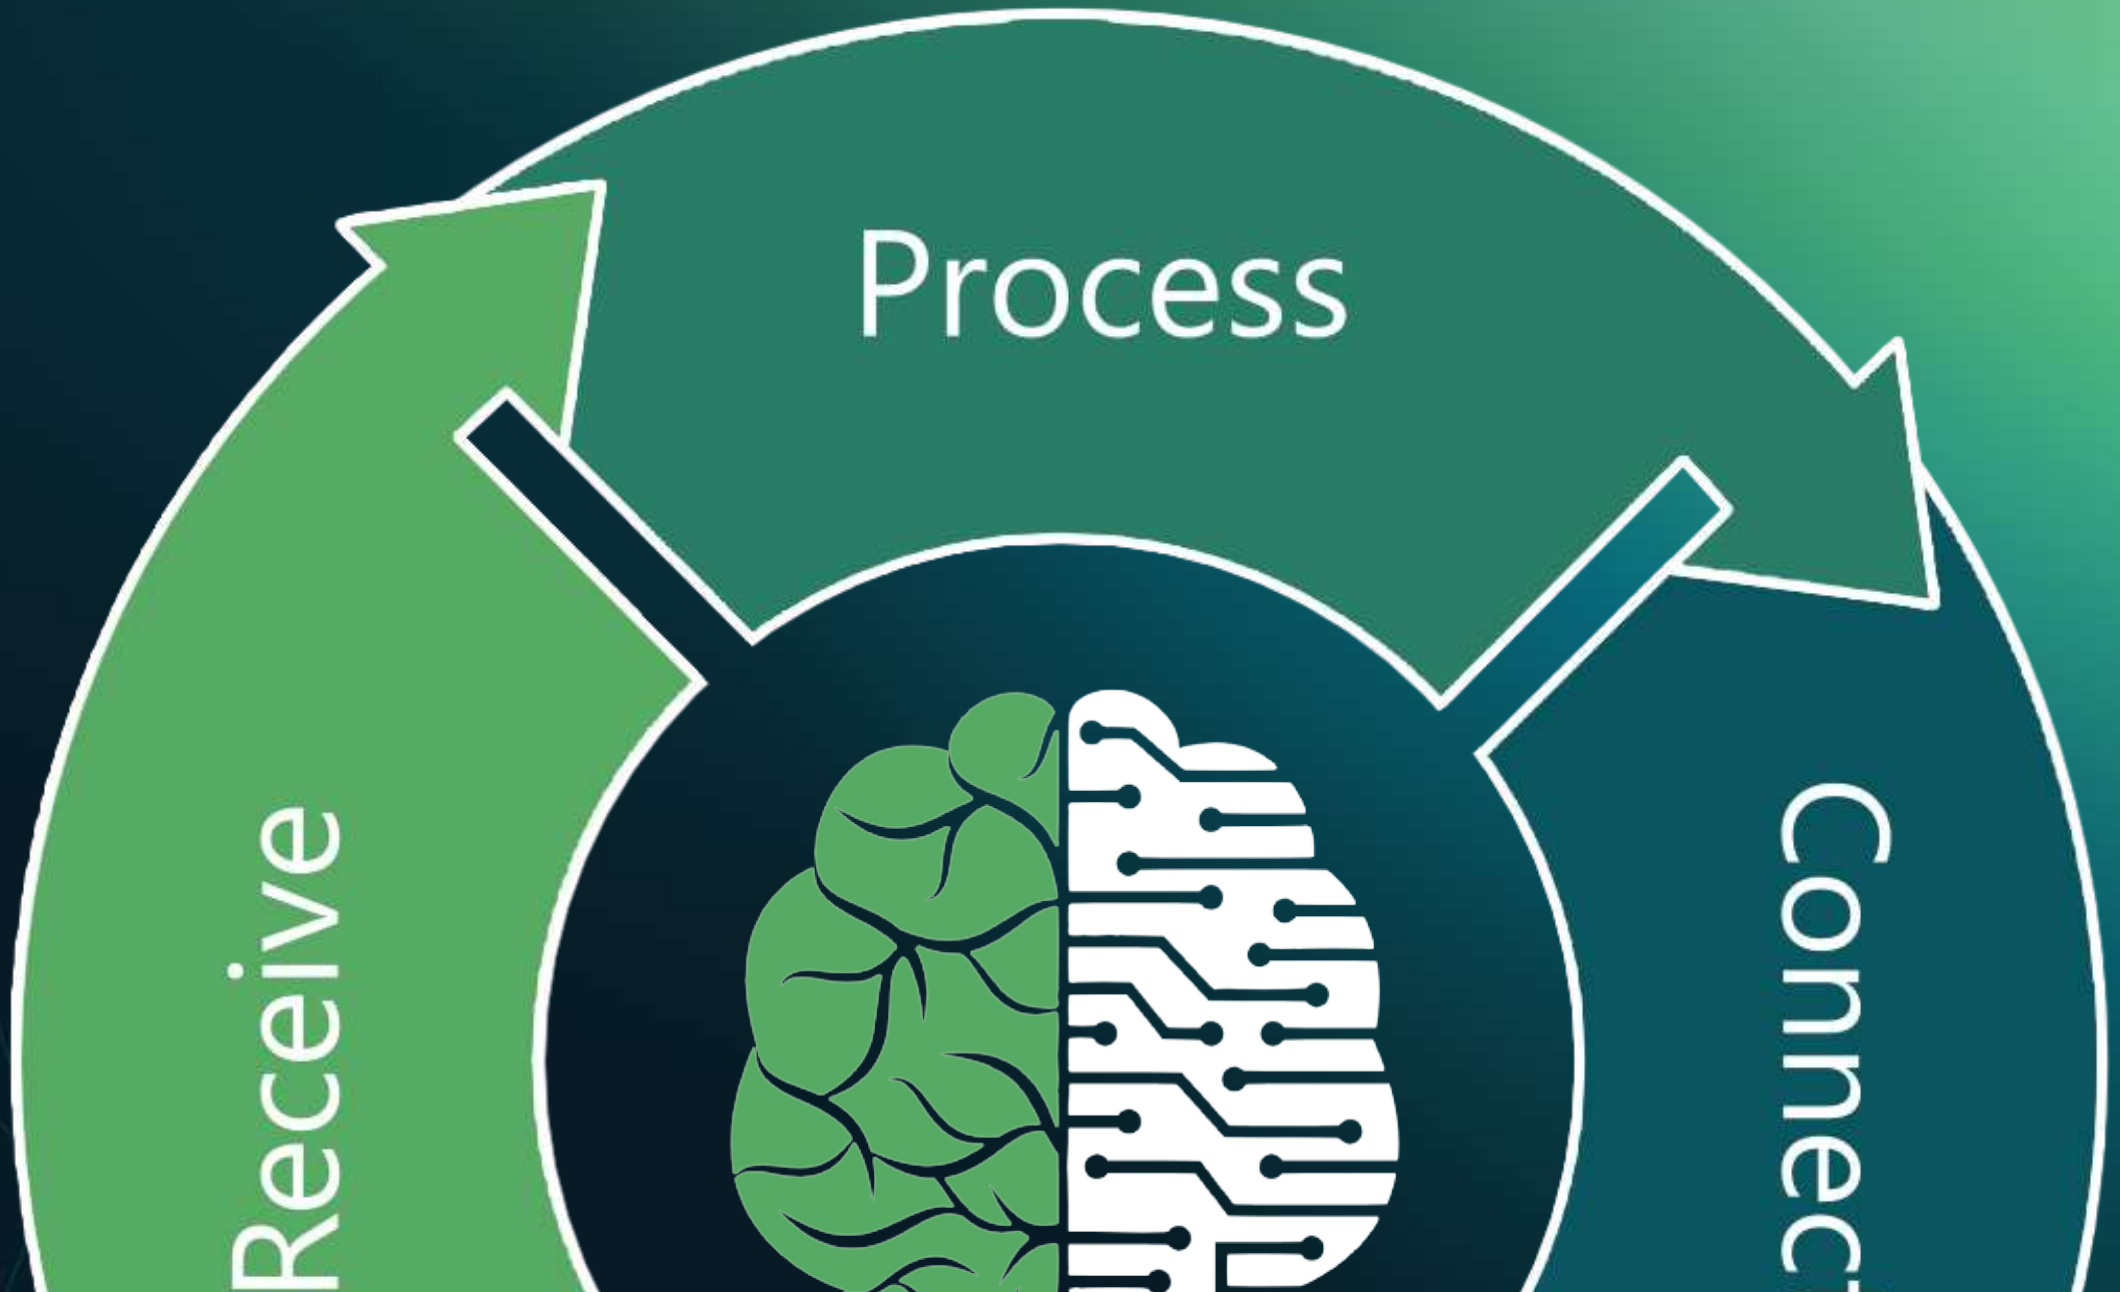
Hello, I am suspecting that the business I work for hands out Order numbers more often than once.
Can you help me write a SQL query which finds all Order Numbers (OrderX) which occur more often than once in the year 2025?

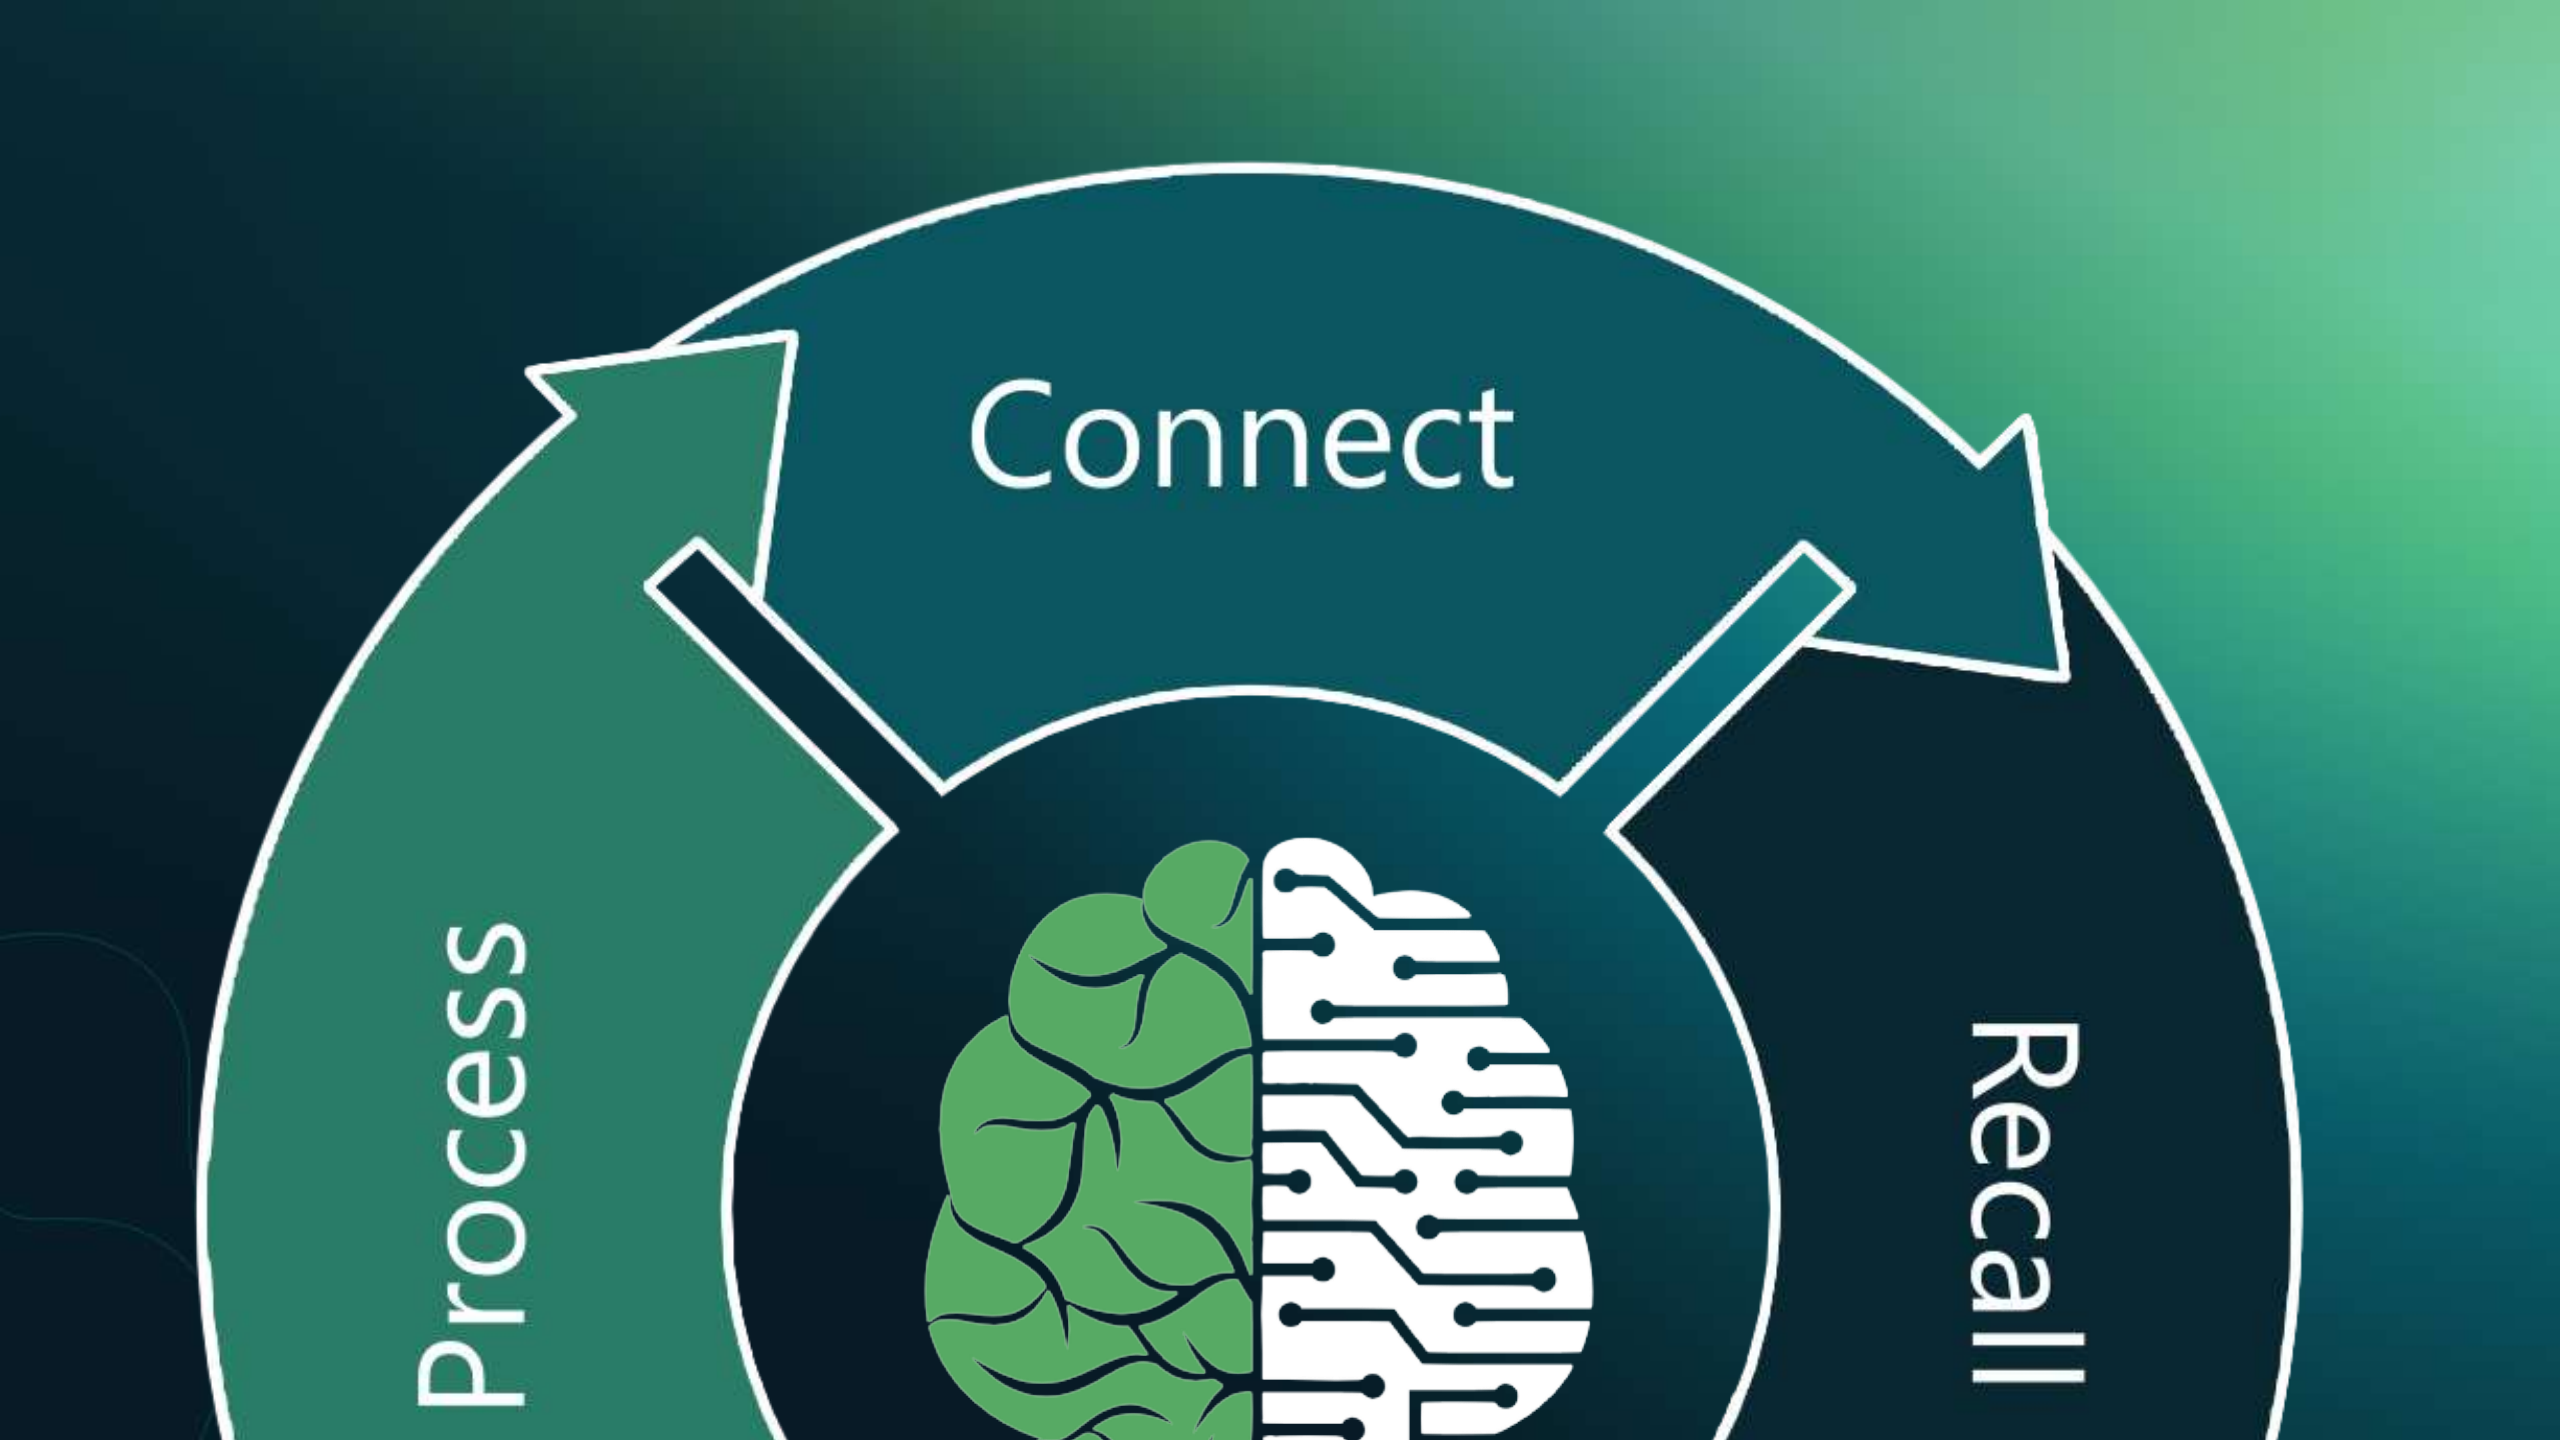
Absolutely! We can write a SQL query that groups by the `OrderX` number, counts how many times each appears, and then filters for those that appear more than once **in the year 2025**. I'll assume your table is

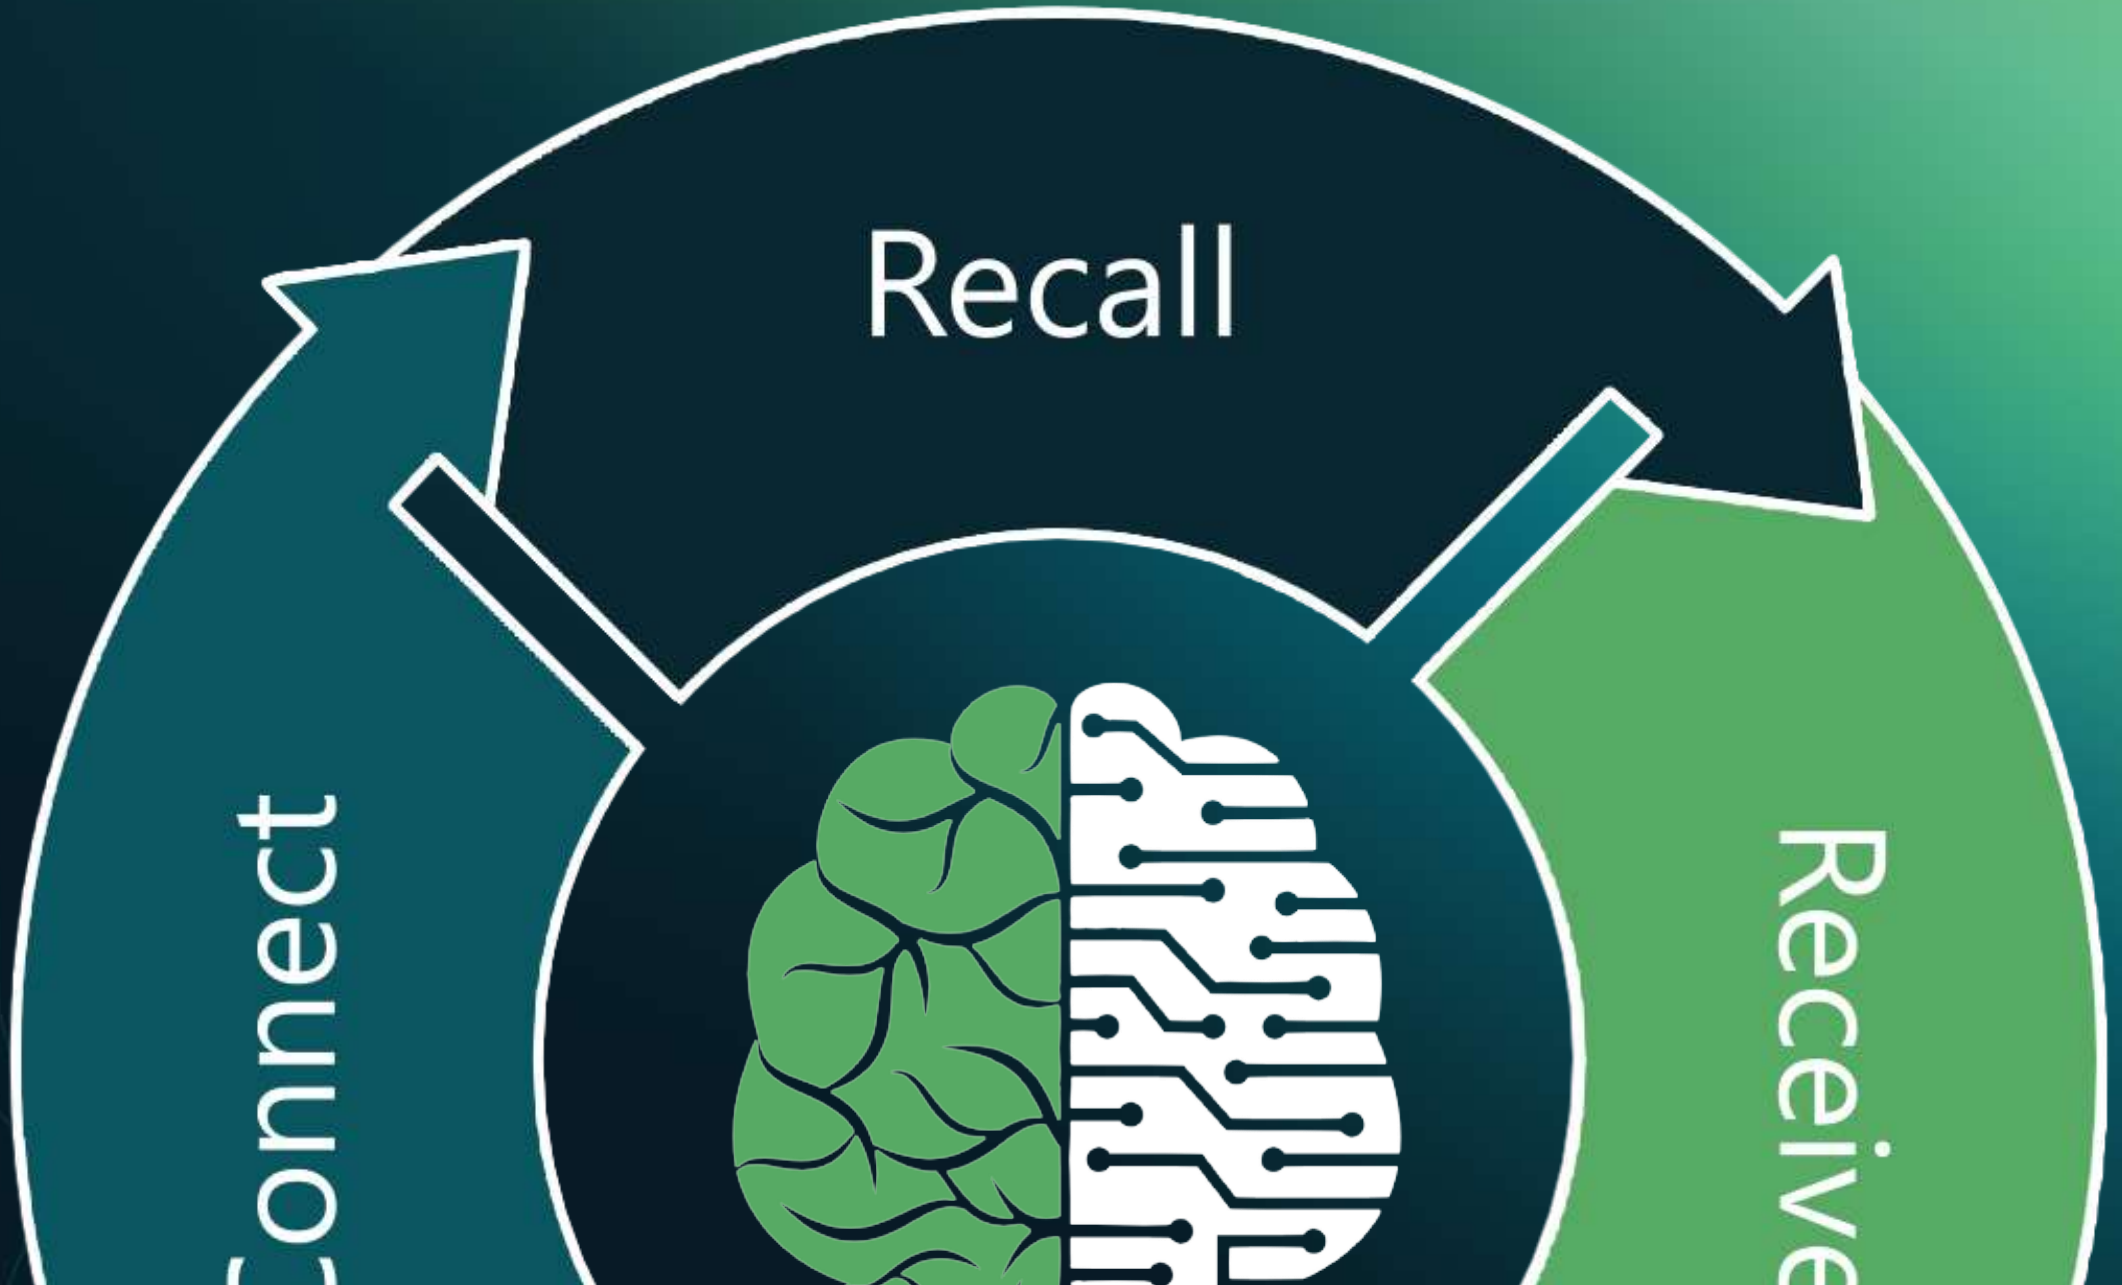
Understanding is NOT retention.

Exposure is NOT expertise.


Receive
Recall
Process


Process
Receive
Connect


Connect
Process
Recall


Recall
Connect
Receive


Receive

Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task△

Nataliya Kosmyna [1]
MIT Media Lab
Cambridge, MA

Eugene Hauptmann
MIT
Cambridge, MA

Ye Tong Yuan
Wellesley College
Wellesley, MA

Jessica Situ
MIT
Cambridge, MA

Xian-Hao Liao
Mass. College of Art and Design (MassArt)
Boston, MA

Ashly Vivian Beresnitzky
MIT
Cambridge, MA

Iris Braunstein
MIT
Cambridge, MA

Pattie Maes
MIT Media Lab
Cambridge, MA

United States

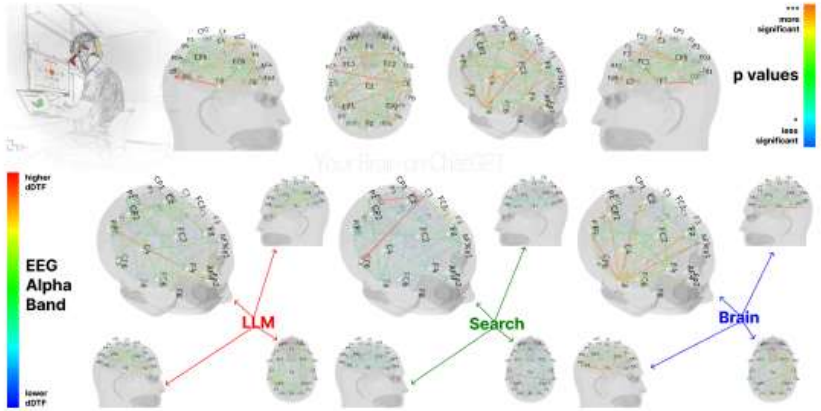


Figure 1. The dynamic Direct Transfer Function (dDTF) EEG analysis of Alpha Band for groups: LLM, Search Engine, Brain-only, including p-values to show significance from moderately significant (*) to highly significant (***).

[1] Nataliya Kosmyna is the corresponding author, please contact her at nkosmyna@mit.edu

△ Distributed under CC BY-NC-SA

Preprint, under review

mit media lab

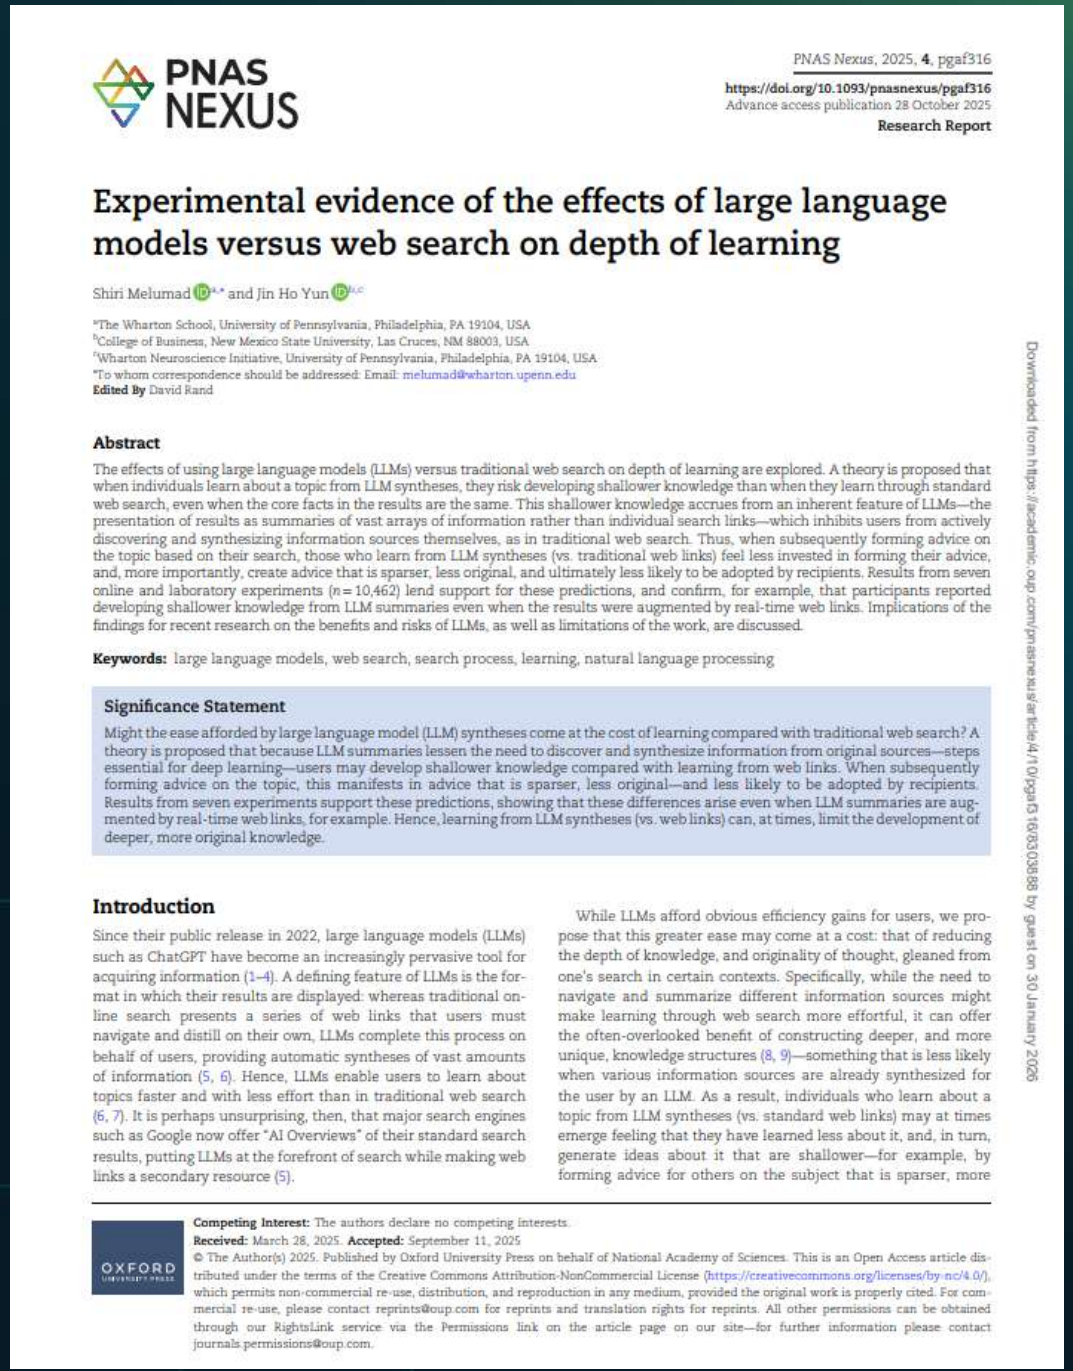
"This study focuses on finding out the **cognitive cost** of using an LLM in the educational context (...)"

"Brain connectivity **systematically scaled down** with the amount of external support: (...)"

"...LLM group fell behind in their **ability to quote** from the essays they wrote just minutes prior."

"We hope this study serves as a preliminary guide to encourage better understanding of the cognitive and practical **impacts of AI on learning environments**."

Scan or Click for Source:



PNAS NEXUS

PNAS Nexus, 2025, 4, pgaf316
https://doi.org/10.1093/pnasnexus/pgaf316
Advance access publication 28 October 2025
Research Report

Experimental evidence of the effects of large language models versus web search on depth of learning

Shiri Melumad [a,*] and Jin Ho Yun [b,c]

[a]The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA
[b]College of Business, New Mexico State University, Las Cruces, NM 88003, USA
[c]Wharton Neuroscience Initiative, University of Pennsylvania, Philadelphia, PA 19104, USA
*To whom correspondence should be addressed: Email: melumad@wharton.upenn.edu
Edited By David Rand

Abstract

The effects of using large language models (LLMs) versus traditional web search on depth of learning are explored. A theory is proposed that when individuals learn about a topic from LLM syntheses, they risk developing shallower knowledge than when they learn through standard web search, even when the core facts in the results are the same. This shallower knowledge accrues from an inherent feature of LLMs—the presentation of results as summaries of vast arrays of information rather than individual search links—which inhibits users from actively discovering and synthesizing information sources themselves, as in traditional web search. Thus, when subsequently forming advice on the topic based on their search, those who learn from LLM syntheses (vs. traditional web links) feel less invested in forming their advice, and, more importantly, create advice that is sparser, less original, and ultimately less likely to be adopted by recipients. Results from seven online and laboratory experiments ($n = 10{,}462$) lend support for these predictions, and confirm, for example, that participants reported developing shallower knowledge from LLM summaries even when the results were augmented by real-time web links. Implications of the findings for recent research on the benefits and risks of LLMs, as well as limitations of the work, are discussed.

Keywords: large language models, web search, search process, learning, natural language processing

Significance Statement

Might the ease afforded by large language model (LLM) syntheses come at the cost of learning compared with traditional web search? A theory is proposed that because LLM summaries lessen the need to discover and synthesize information from original sources—steps essential for deep learning—users may develop shallower knowledge compared with learning from web links. When subsequently forming advice on the topic, this manifests in advice that is sparser, less original—and less likely to be adopted by recipients. Results from seven experiments support these predictions, showing that these differences arise even when LLM summaries are augmented by real-time web links, for example. Hence, learning from LLM syntheses (vs. web links) can, at times, limit the development of deeper, more original knowledge.

Introduction

Since their public release in 2022, large language models (LLMs) such as ChatGPT have become an increasingly pervasive tool for acquiring information (1–4). A defining feature of LLMs is the format in which their results are displayed: whereas traditional online search presents a series of web links that users must navigate and distill on their own, LLMs complete this process on behalf of users, providing automatic syntheses of vast amounts of information (5, 6). Hence, LLMs enable users to learn about topics faster and with less effort than in traditional web search (6, 7). It is perhaps unsurprising, then, that major search engines such as Google now offer "AI Overviews" of their standard search results, putting LLMs at the forefront of search while making web links a secondary resource (5).

While LLMs afford obvious efficiency gains for users, we propose that this greater ease may come at a cost: that of reducing the depth of knowledge, and originality of thought, gleaned from one's search in certain contexts. Specifically, while the need to navigate and summarize different information sources might make learning through web search more effortful, it can offer the often-overlooked benefit of constructing deeper, and more unique, knowledge structures (8, 9)—something that is less likely when various information sources are already synthesized for the user by an LLM. As a result, individuals who learn about a topic from LLM syntheses (vs. standard web links) may at times emerge feeling that they have learned less about it, and, in turn, generate ideas about it that are shallower—for example, by forming advice for others on the subject that is sparser, more

Competing Interest: The authors declare no competing interests.
Received: March 28, 2025. **Accepted:** September 11, 2025
© The Author(s) 2025. Published by Oxford University Press on behalf of National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (https://creativecommons.org/licenses/by-nc/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

Downloaded from https://academic.oup.com/pnasnexus/article/4/10/pgaf316/8303688 by guest on 30 January 2026

"The effects of using LLMs versus traditional web search on depth of learning are explored."

"(…) shallower knowledge accrues from an inherent feature of LLMs – the presentation of results as summaries of vast arrays of information rather than individual search links – (…)"

"(…) those who learn from LLM syntheses feel less invested in forming their advice, and, more importantly, create advice that is sparser, less original, and ultimately less likely to be adopted by recipients."

Scan or Click for Source:



The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers

Hao-Ping (Hank) Lee
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
haopingl@cs.cmu.edu

Advait Sarkar
Microsoft Research
Cambridge, United Kingdom
advait@microsoft.com

Lev Tankelevitch
Microsoft Research
Cambridge, United Kingdom
levt@microsoft.com

Ian Drosos
Microsoft Research
Cambridge, United Kingdom
t-iandrosos@microsoft.com

Sean Rintel
Microsoft Research
Cambridge, United Kingdom
serintel@microsoft.com

Richard Banks
Microsoft Research Cambridge
Cambridge, United Kingdom
rbanks@microsoft.com

Nicholas Wilson
Microsoft Research
Cambridge, United Kingdom
niwilson@microsoft.com

Abstract

The rise of Generative AI (GenAI) in knowledge workflows raises questions about its impact on critical thinking skills and practices. We survey 319 knowledge workers to investigate 1) when and how they perceive the enaction of critical thinking when using GenAI, and 2) when and why GenAI affects their effort to do so. Participants shared 936 first-hand examples of using GenAI in work tasks. Quantitatively, when considering both task- and user-specific factors, a user's task-specific self-confidence and confidence in GenAI are predictive of whether critical thinking is enacted and the effort of doing so in GenAI-assisted tasks. Specifically, higher confidence in GenAI is associated with less critical thinking, while higher self-confidence is associated with more critical thinking. Qualitatively, GenAI shifts the nature of critical thinking toward information verification, response integration, and task stewardship. Our insights reveal new design challenges and opportunities for developing GenAI tools for knowledge work.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**.

Keywords

Critical thinking, Generative AI tools, Knowledge worker, Bloom's taxonomy, Survey

ACM Reference Format:

Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 22 pages. https://doi.org/10.1145/3706598.3713778

This work is licensed under a Creative Commons Attribution 4.0 International License.

CHI '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1394-1/25/04
https://doi.org/10.1145/3706598.3713778

1 Introduction

Generative AI (**GenAI**) tools, defined as any *"end user tool [...] whose technical implementation includes a generative model based on deep learning"*,[1] are the latest in a long line of technologies that raise questions about their impact on the quality of human thought, a line that includes writing (objected to by Socrates), printing (objected to by Trithemius), calculators (objected to by teachers of arithmetic), and the Internet.

Such consternation is not unfounded. Used improperly, technologies can and do result in the deterioration of cognitive faculties that ought to be preserved. As Bainbridge [7] noted, a key irony of automation is that by mechanizing routine tasks and leaving exception-handling to the human user, you deprive the user of the routine opportunities to practice their judgment and strengthen their cognitive musculature, leaving them atrophied and unprepared when the exceptions do arise.

In response, research has begun looking closely at how different activities are impacted by GenAI and the extent to which cognitive offloading [8] occurs, and whether this may be an undesirable thing. Some work has focused, for instance, on studying the effects of GenAI use on memory (e.g., [1, 106]) and on creativity (e.g., [28, 100]). Moreover, design research has also been developing interventions that *improve* the ability of people to think in certain ways (e.g., [24]). We review these lines of work in Section 2.

In this paper, we focus on a higher-level concept that captures another aspect of thought considered desirable and worthy of preservation: *critical thinking* (defined in Section 2). The effect of the use

[1] While there is no broad consensus on how to define this now-common term, for clarity we adopt this definition, a rationale for which is given in [115].

Microsoft Research

"GenAI tools are the latest in a long line of technologies that raise questions about their impact on the quality of human thought (...)"

"Output diversity is a proxy for critical thinking, and a flawed one."

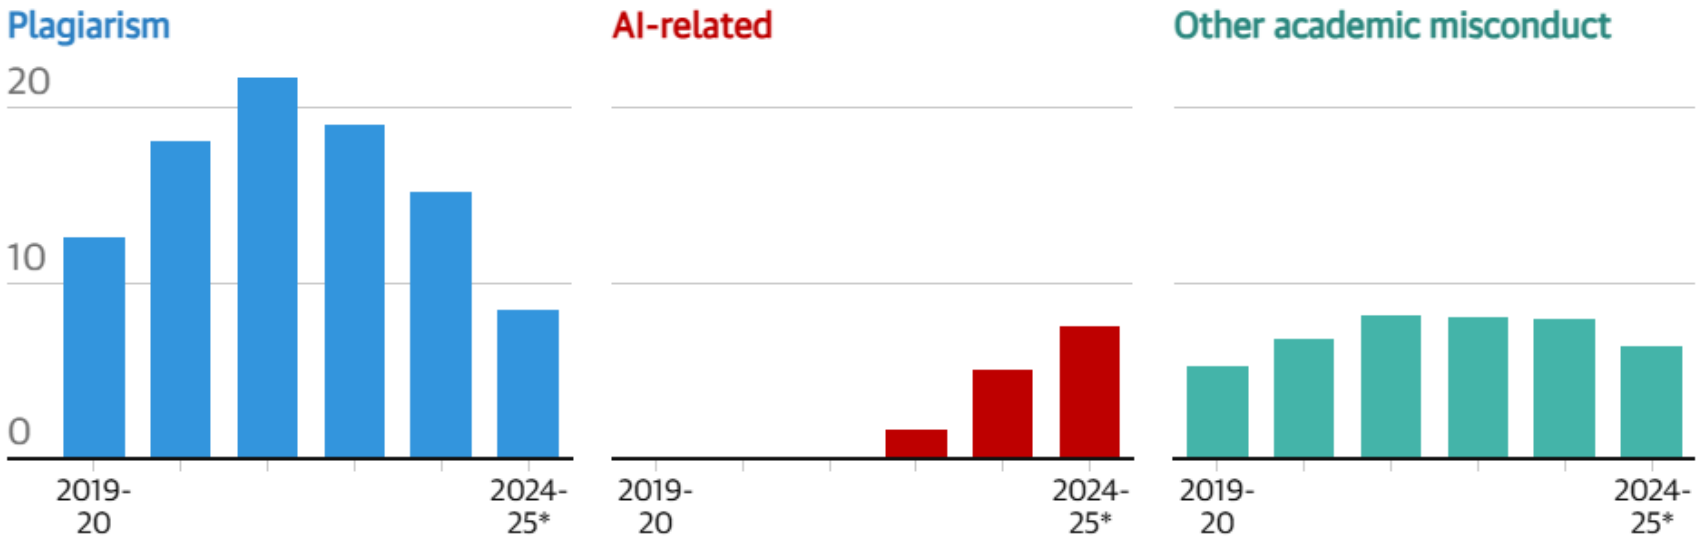
"We find that GenAI tools reduce the perceived effort of critical thinking while also encouraging over-reliance on AI, with confidence in the tool often diminishing independent problem solving."

"(...) workers who are confident in their own skills tend to perceive greater effort in (critical thinking) tasks, particulary when evaluating and applying AI responses."

Scan or Click for Source:

As plagiarism falls, AI-related misconduct is rising

Proven misconduct cases per 1,000 students



Guardian graphic. Source: Guardian FoI requests . Notes: Between 80 and 125 universities responded depending on year and misconduct category. Rates are per 1,000 student
projected based on data to May.

% of videos shown to new
s are 'AI slop', study finds

ed content is now saturating social
about $117m a year, data shows



Revealed: Thousands of UK university students caught cheating using AI

Guardian investigation finds almost 7,000 proven cases of cheating – and experts says these are tip of the iceberg

Your Smart fash

Build your digital wardrobe and get AI-powered outfit r

Manage Wardrobe

02:53pm EST

Why this matters for Data Literacy

- Some **awareness**

- **Convenient tool** becomes a **convenient excuse**
    - Why train knowledge that everybody already has available on prompt?
    - Why motivate yourself to learn something when it is so easy to machinate?

- Ironically: AI will **raise the standard** for everyone, not the other way around
    - If AI is already providing a high amount of value, the **threshold of expertise** one needs to contribute will become higher than ever

How to develop **real expertise** in yourself and your organization – so that when you walk into a room, you know that any data, tool or problem you interact with, **you can produce a higher quality result than anyone else.**



THANK YOU!

