



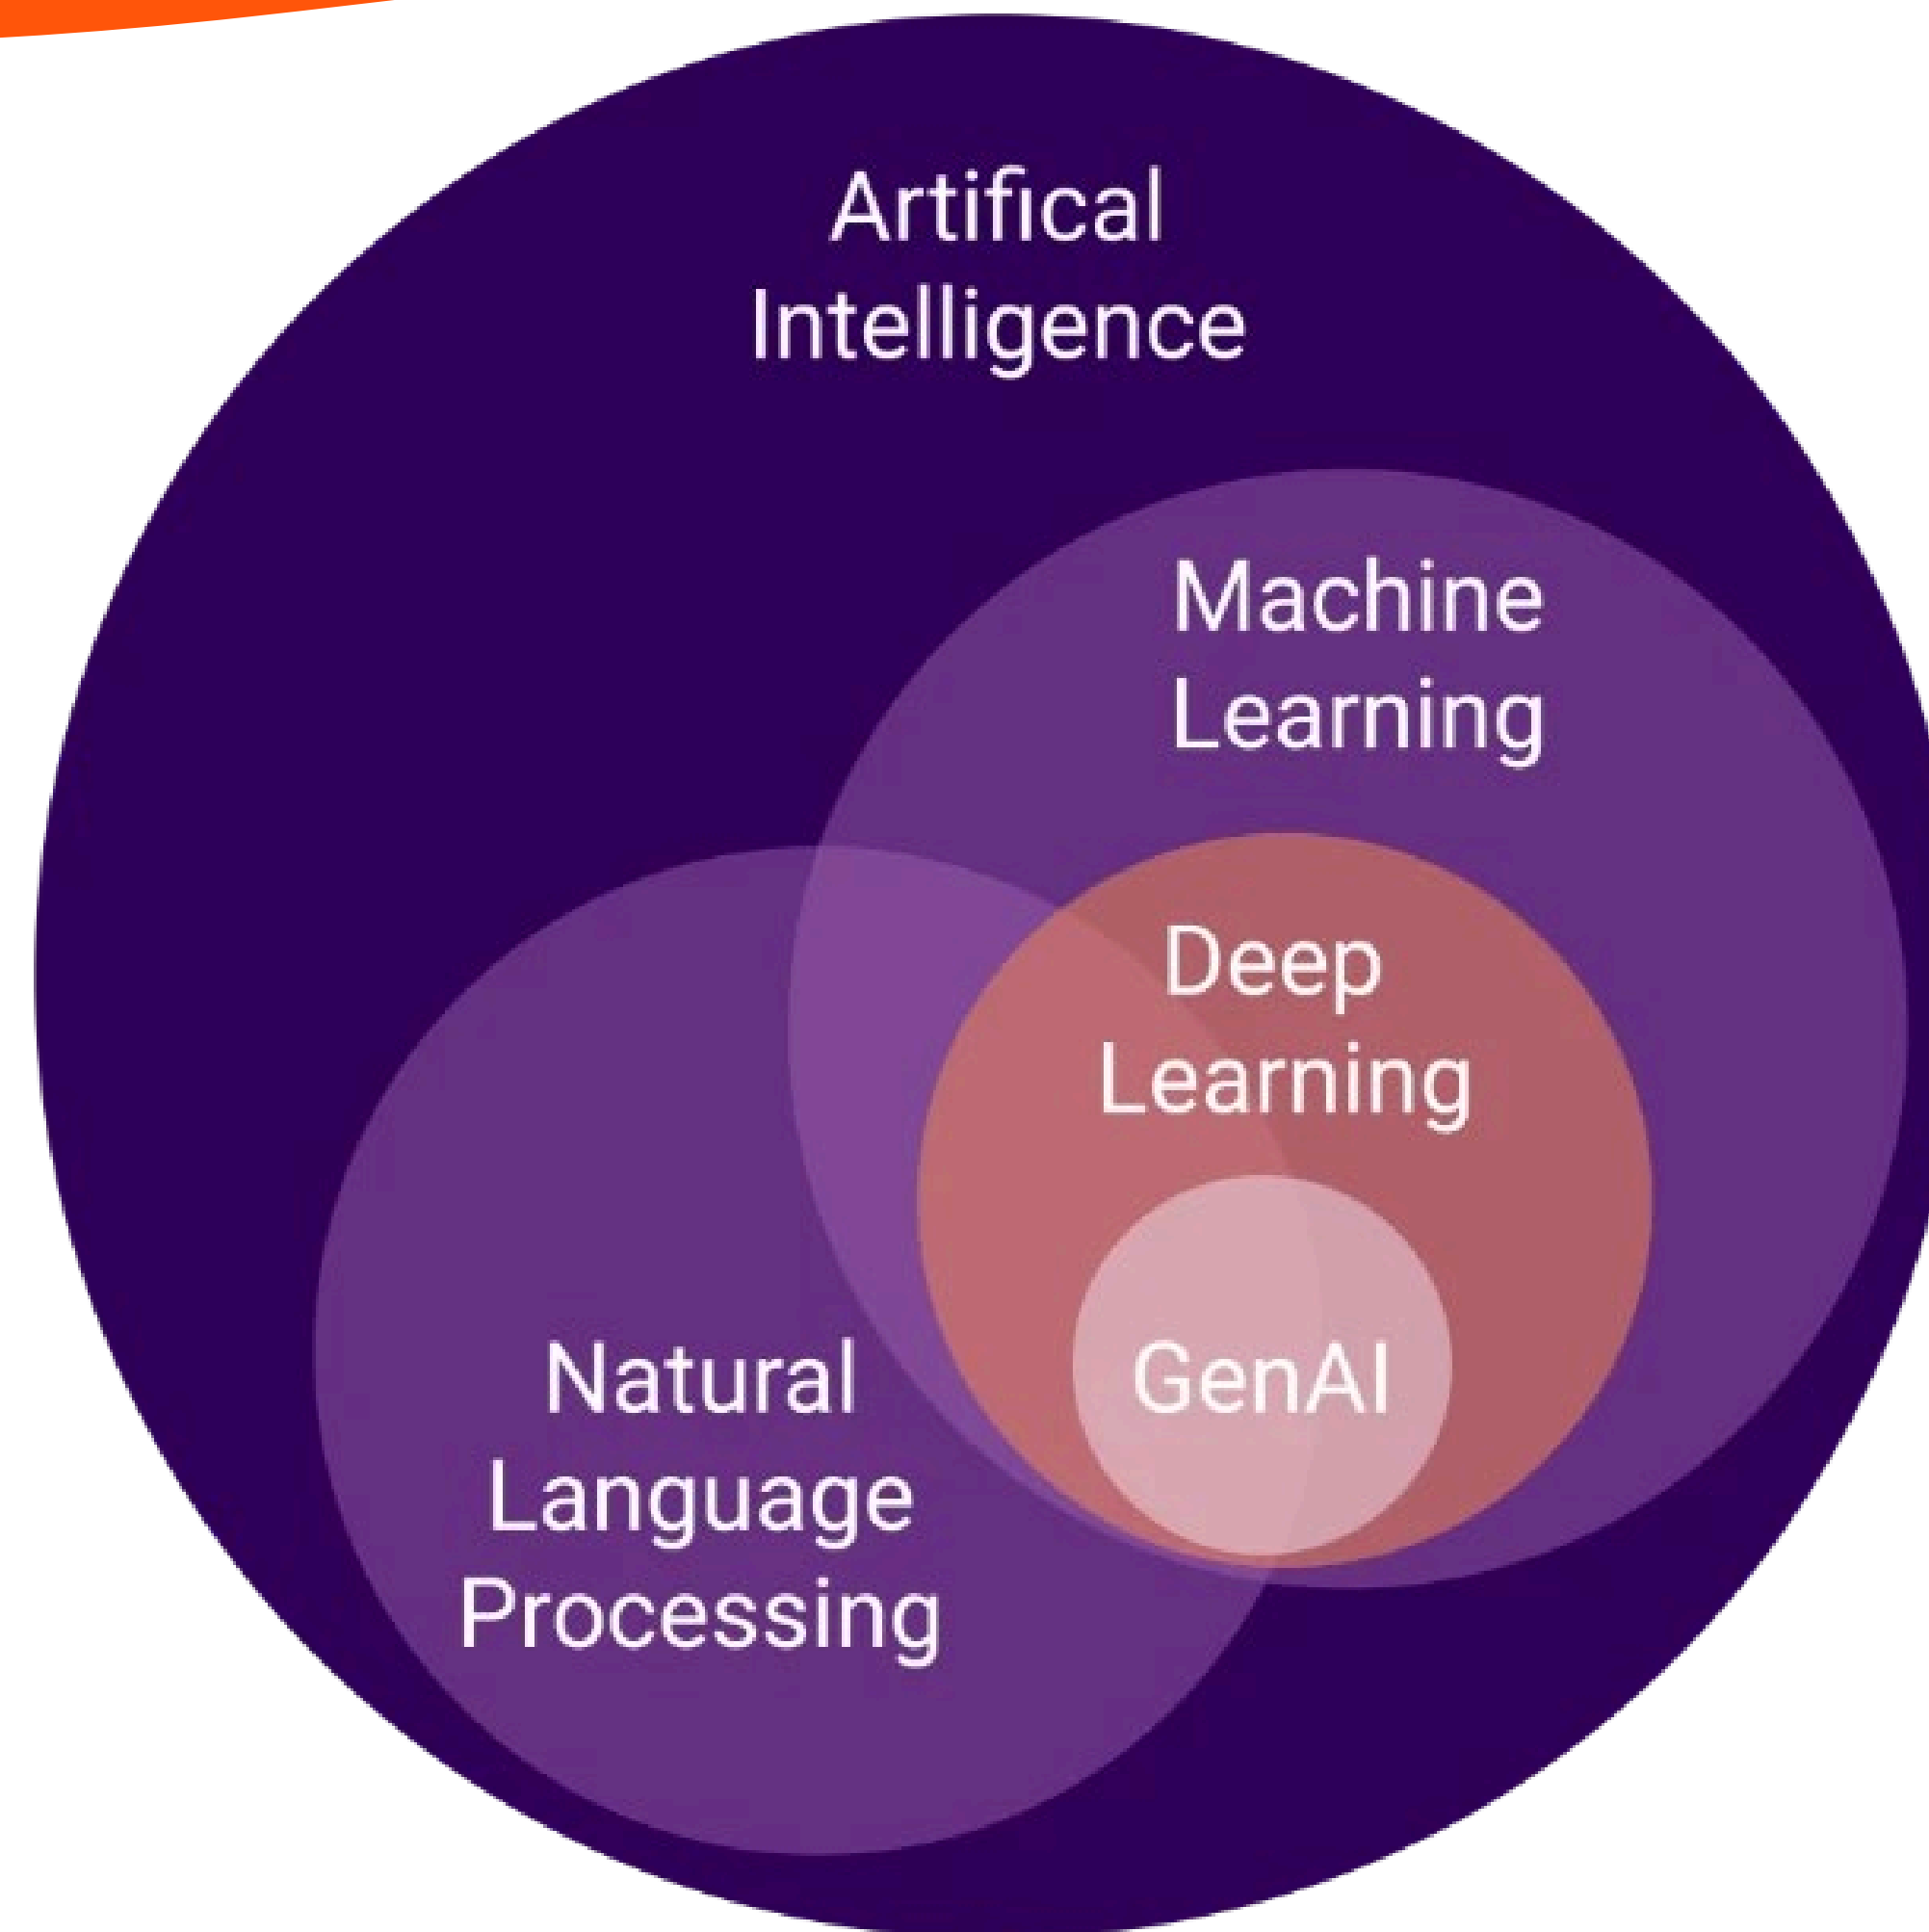
AI IN DATAFLEX

KAI.CHEN

WHAT IS AI

WHAT IS AI

AI IS A HUGE UMBRELLA



WHAT IS AI

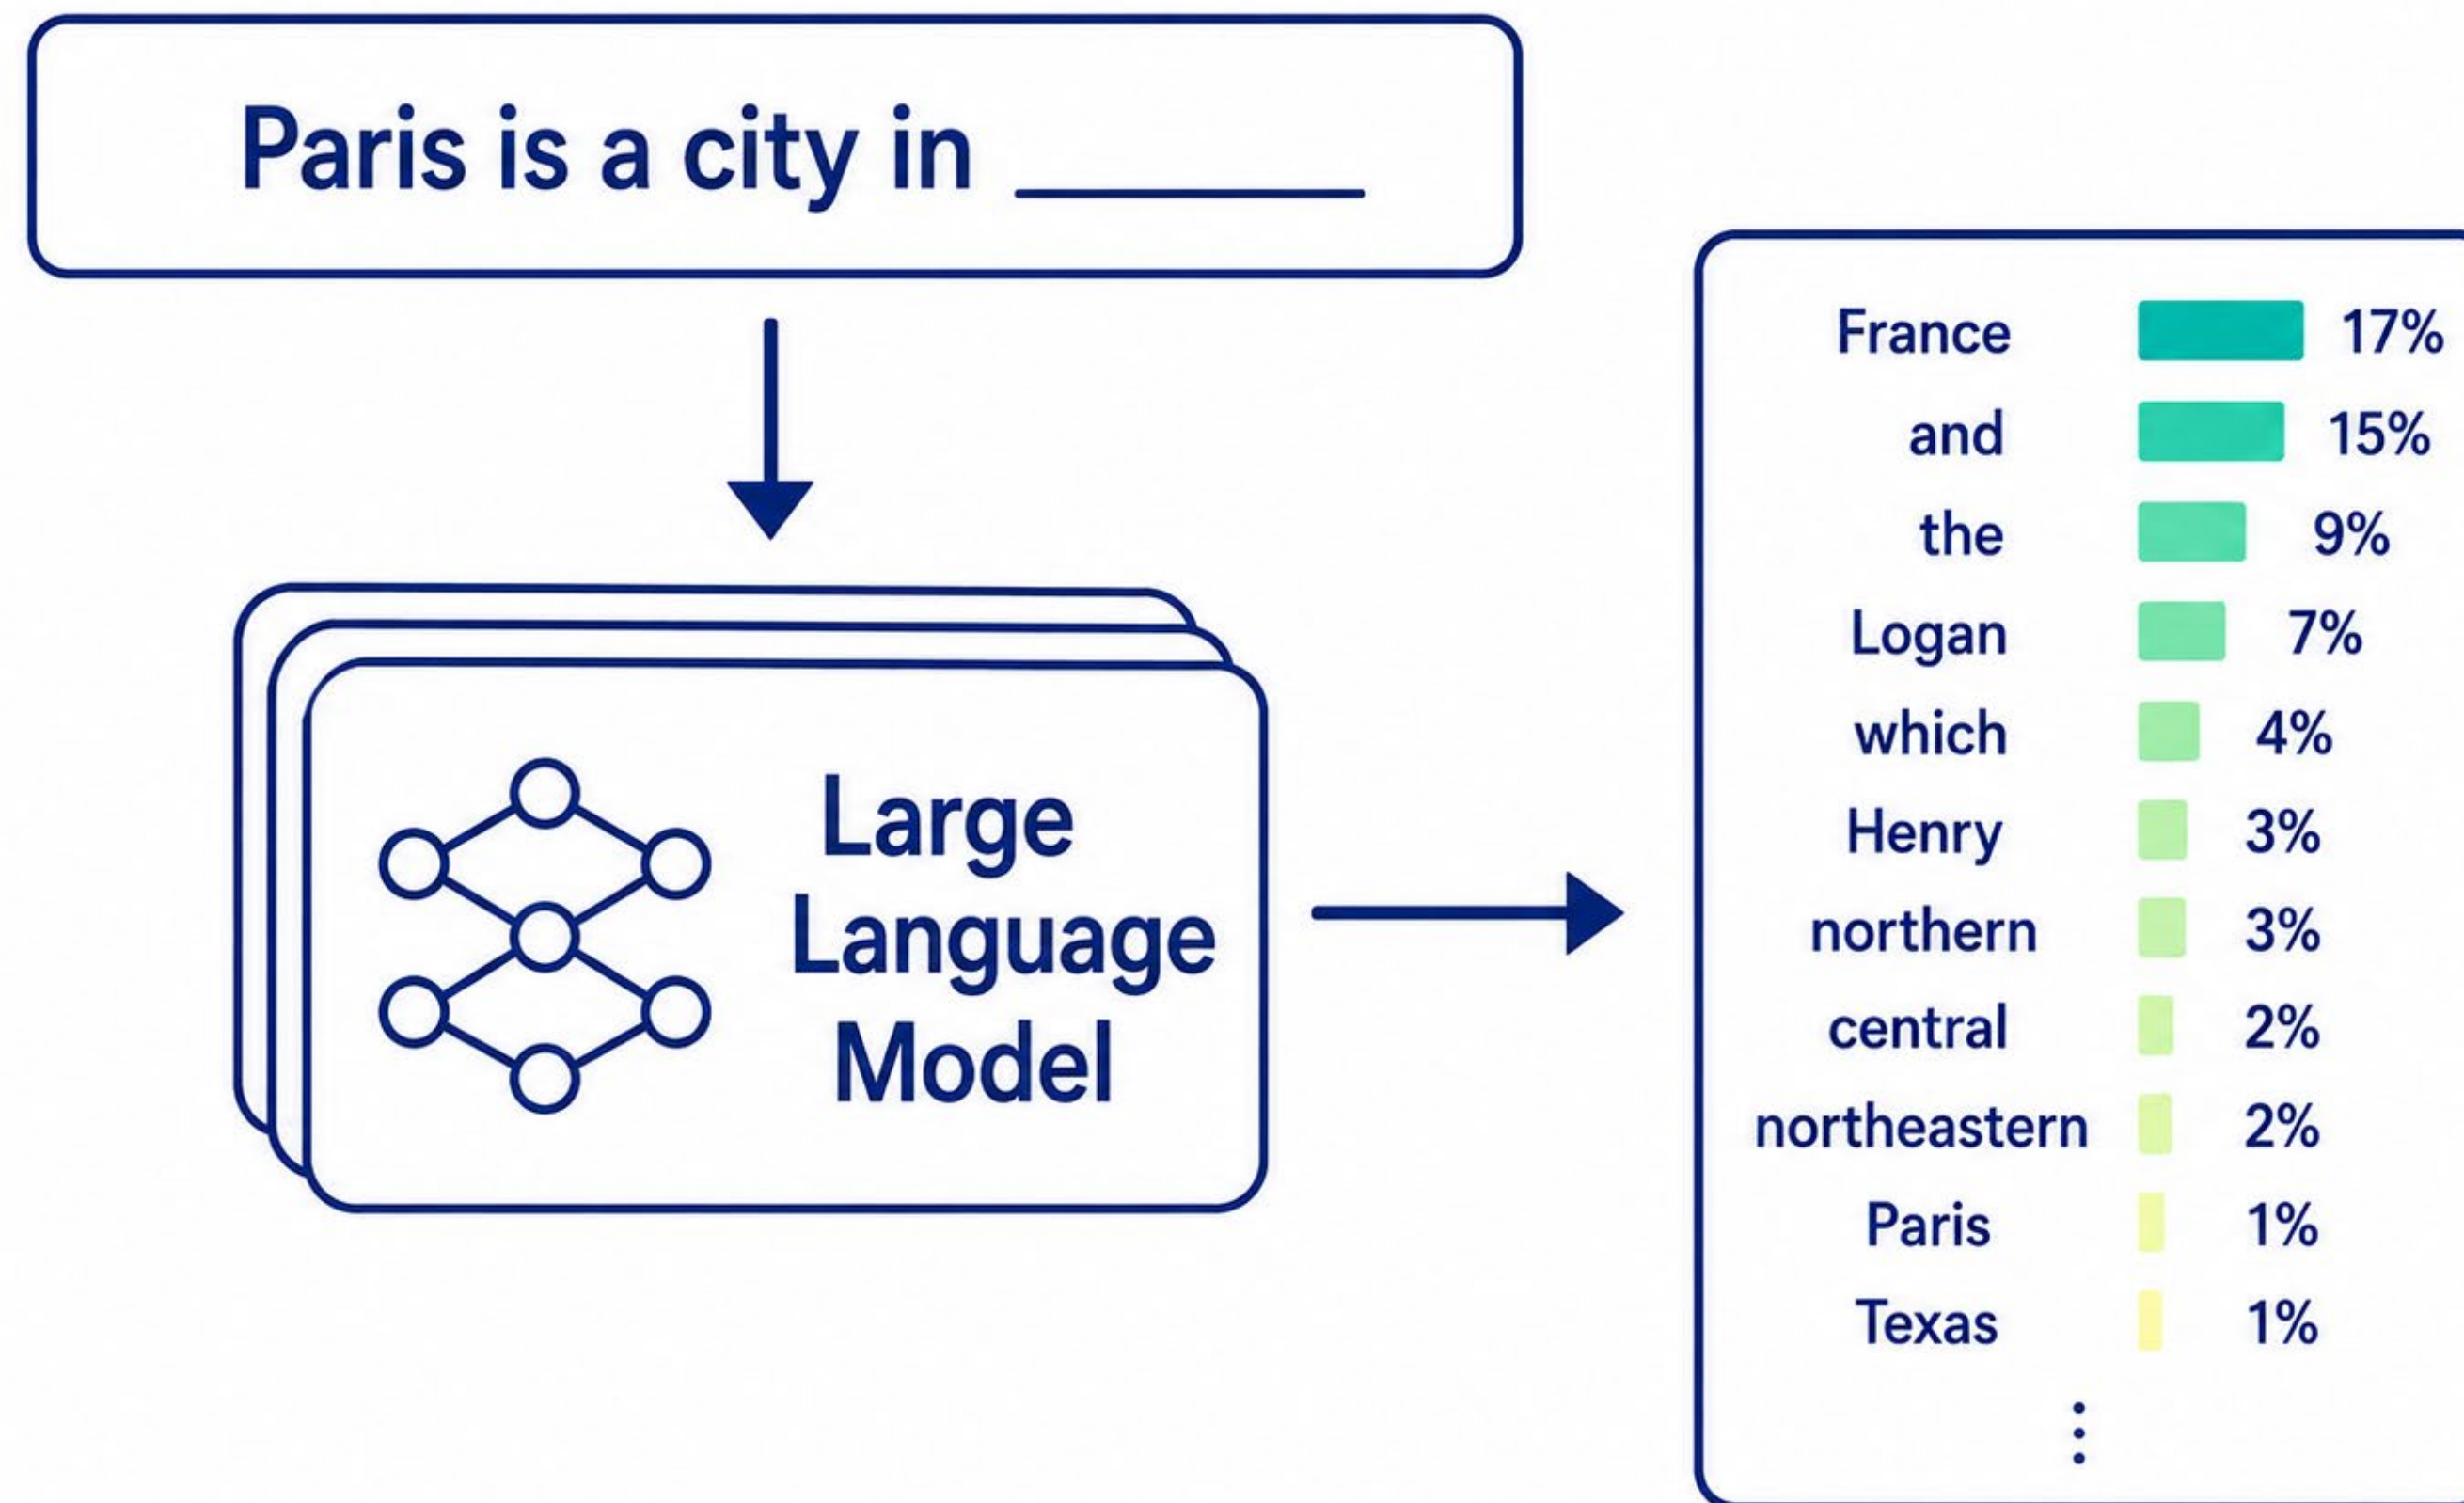
WHAT I'LL SHOW YOU TODAY

1. How an LLM actually works
2. RAG – giving the AI reference documents
3. Demo: the DataFlex docs AI chat
4. MCP – giving the AI hands
5. Demo: AI in your SQL databases
6. Demo: the AI Agent inside the Studio

HOW AN LLM WORKS

HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

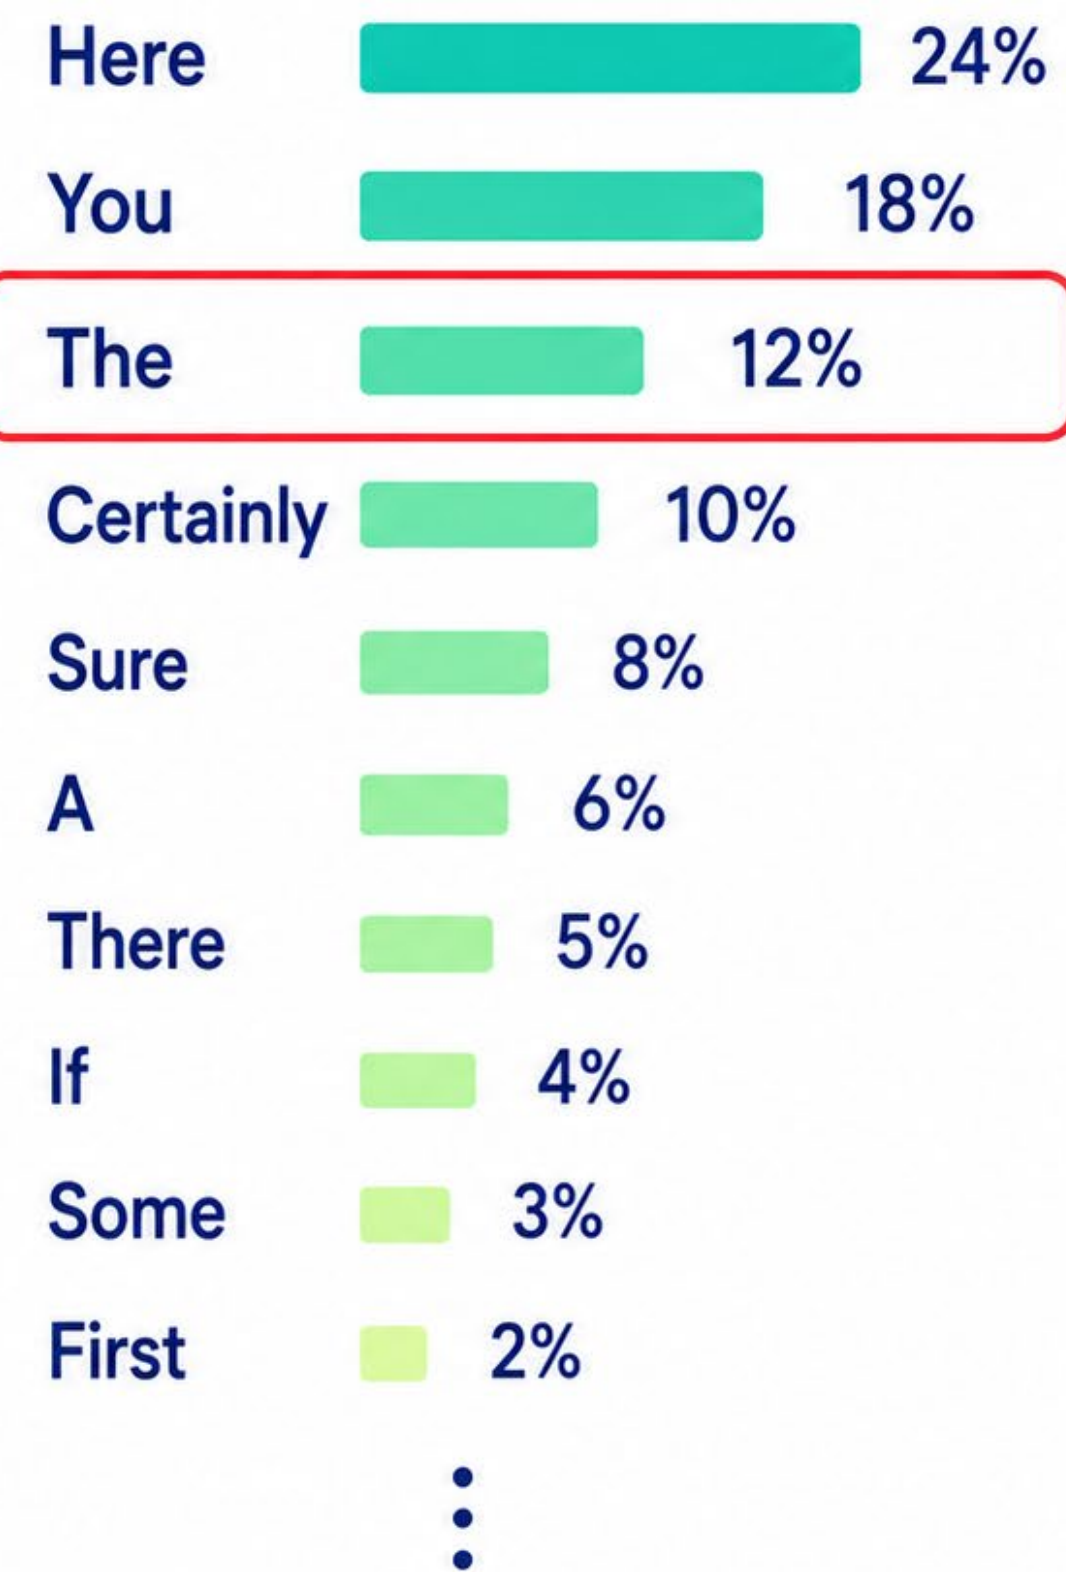


HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: _____

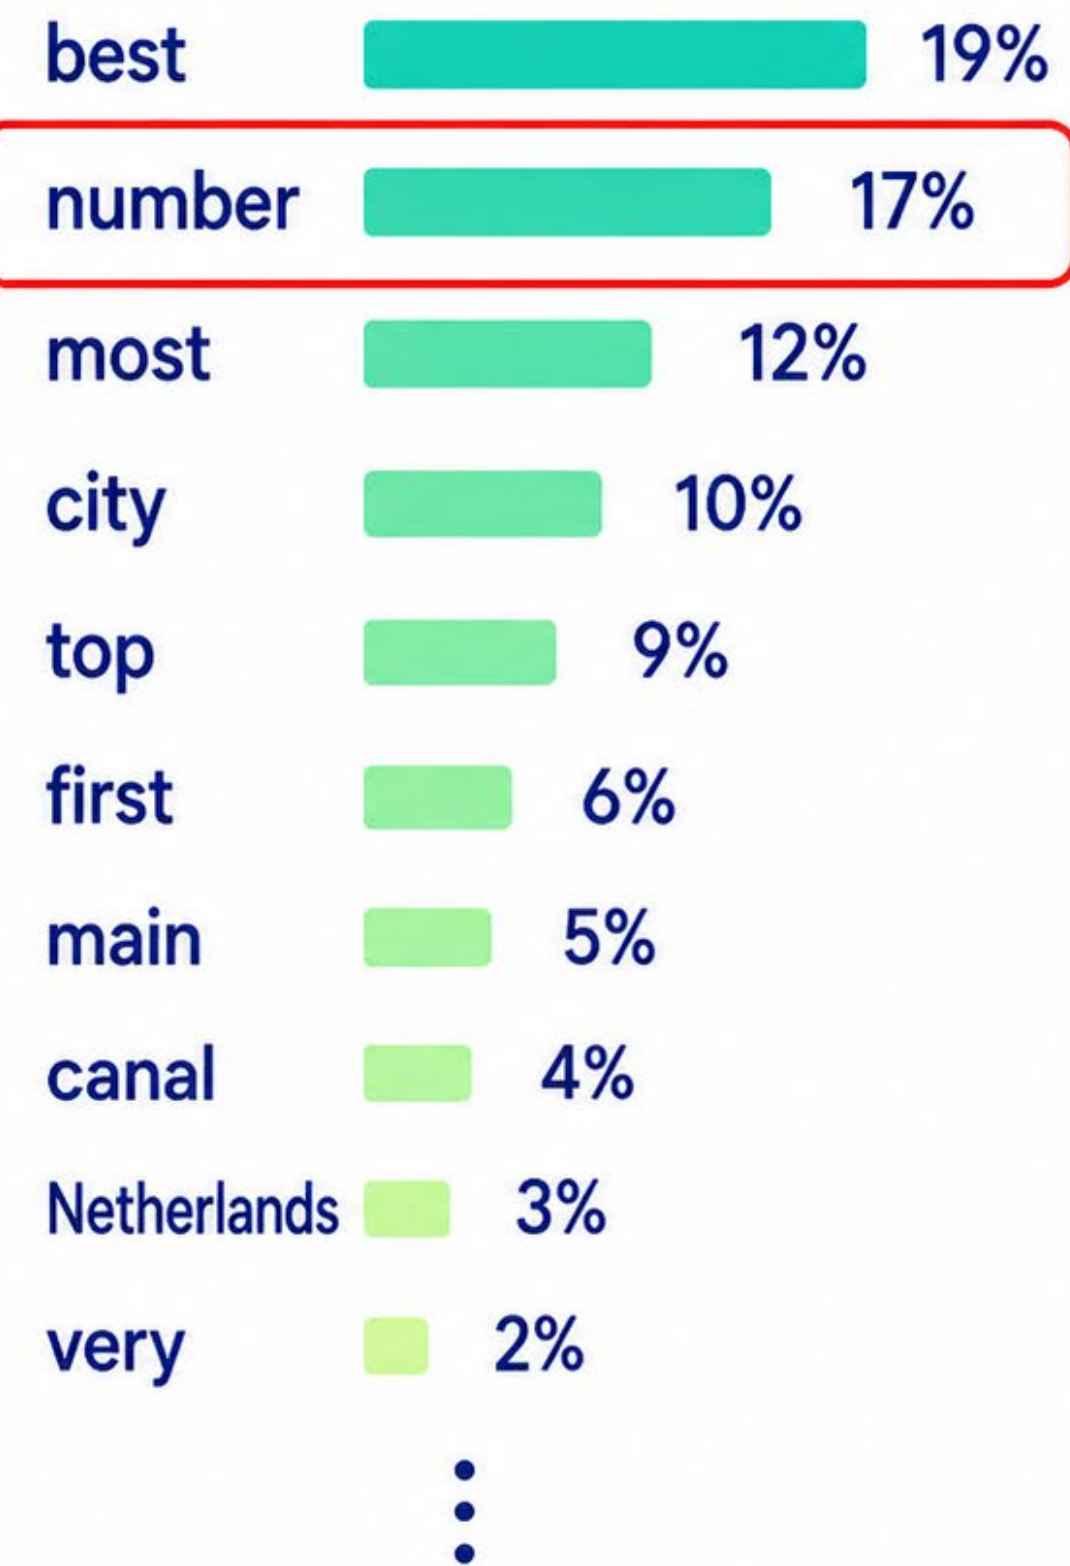


HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The _____



HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number ___

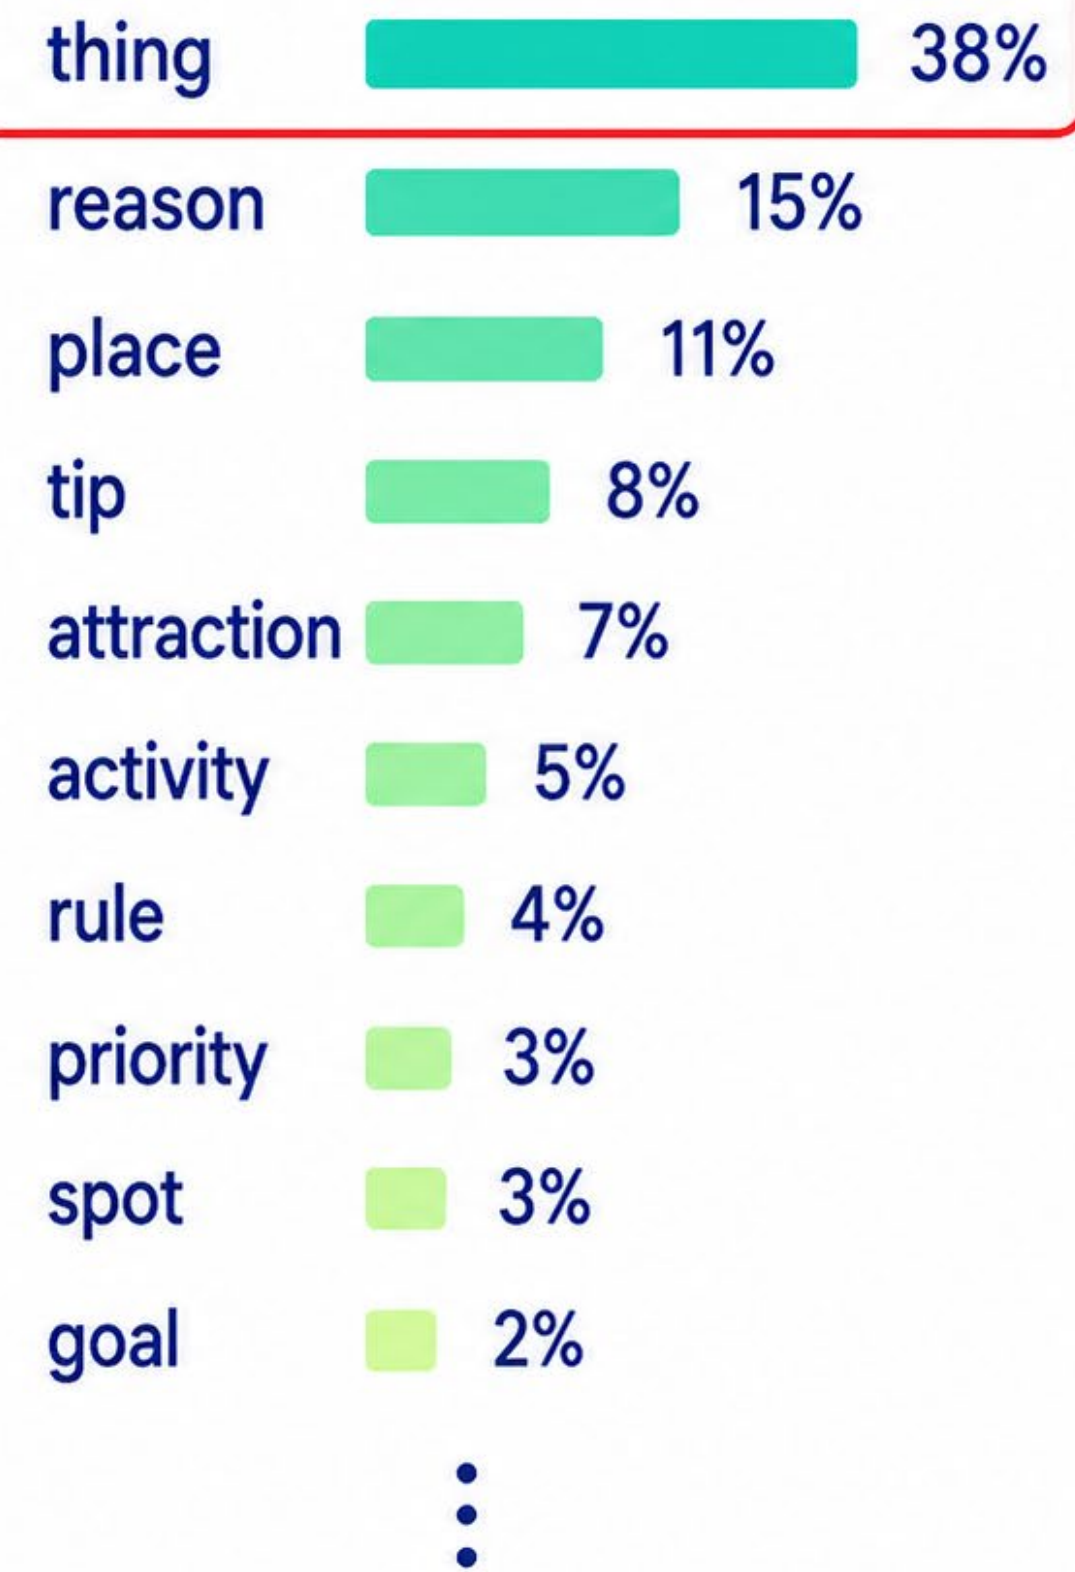


HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number 1 ___



HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number 1 thing ____

to 46%

you 9%

for 8%

about 6%

I 5%

that 4%

people 3%

when 3%

visitors 2%

everyone 2%

⋮

HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number 1 thing to ___

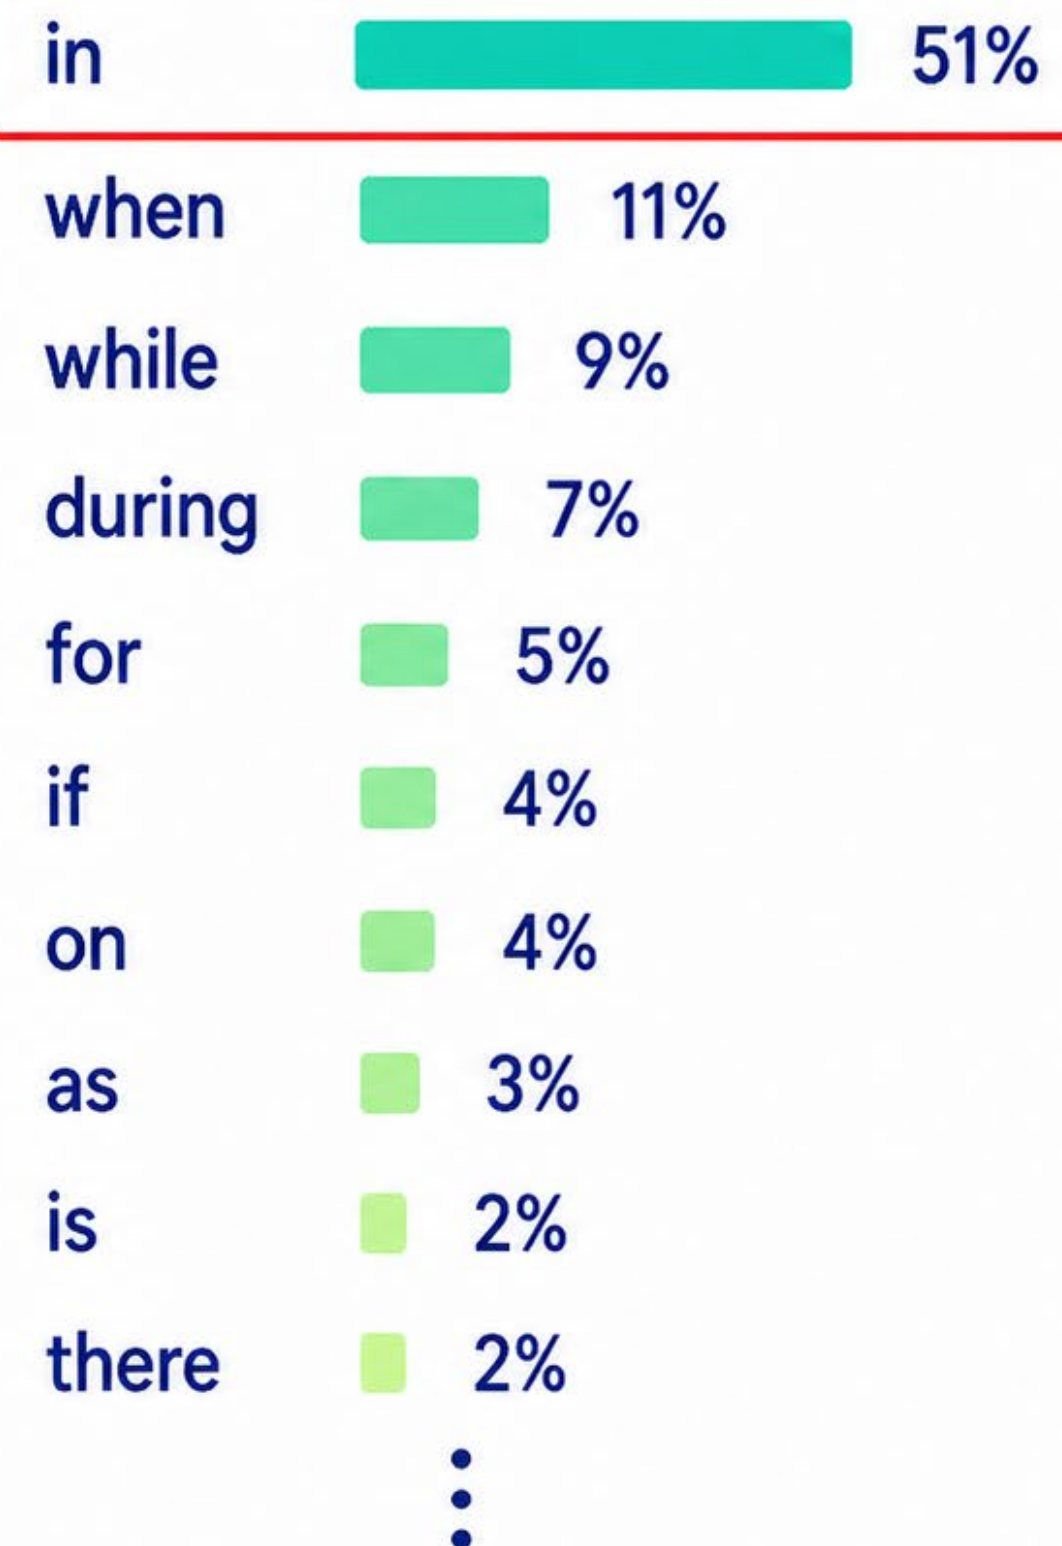


HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number 1 thing to do __



HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number 1 thing to do in Amsterdam ____



HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number 1 thing to do in Amsterdam is to attend _

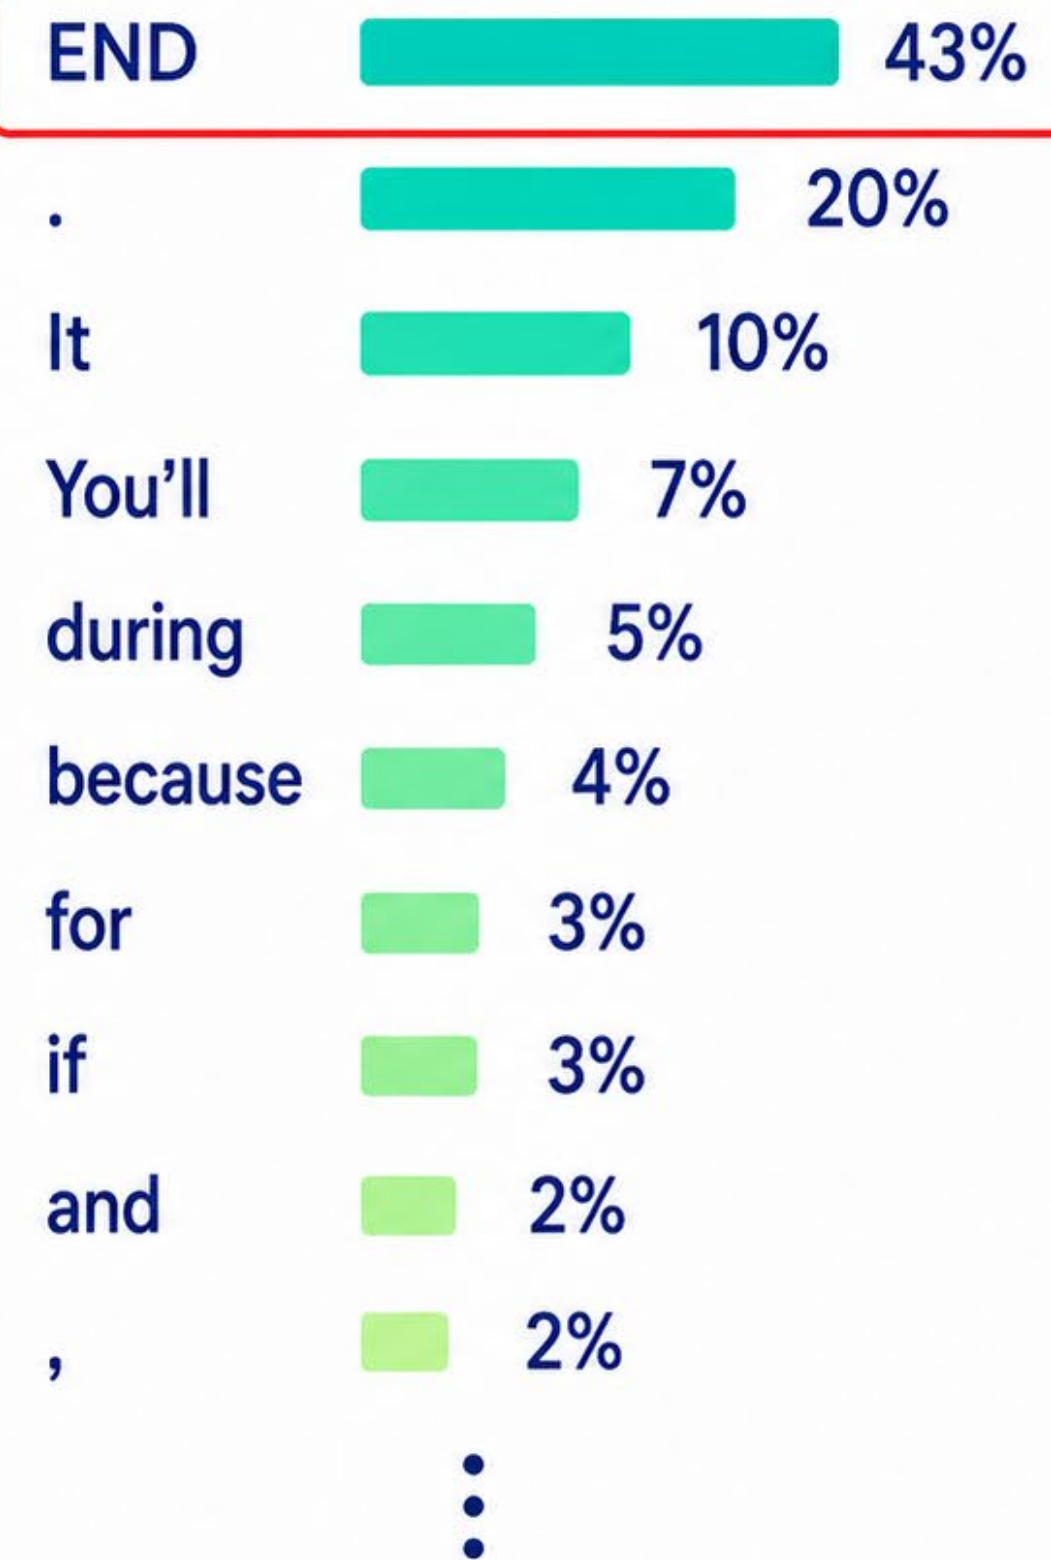


HOW AN LLM WORKS

“IT’S AUTOCOMPLETE. REALLY.”

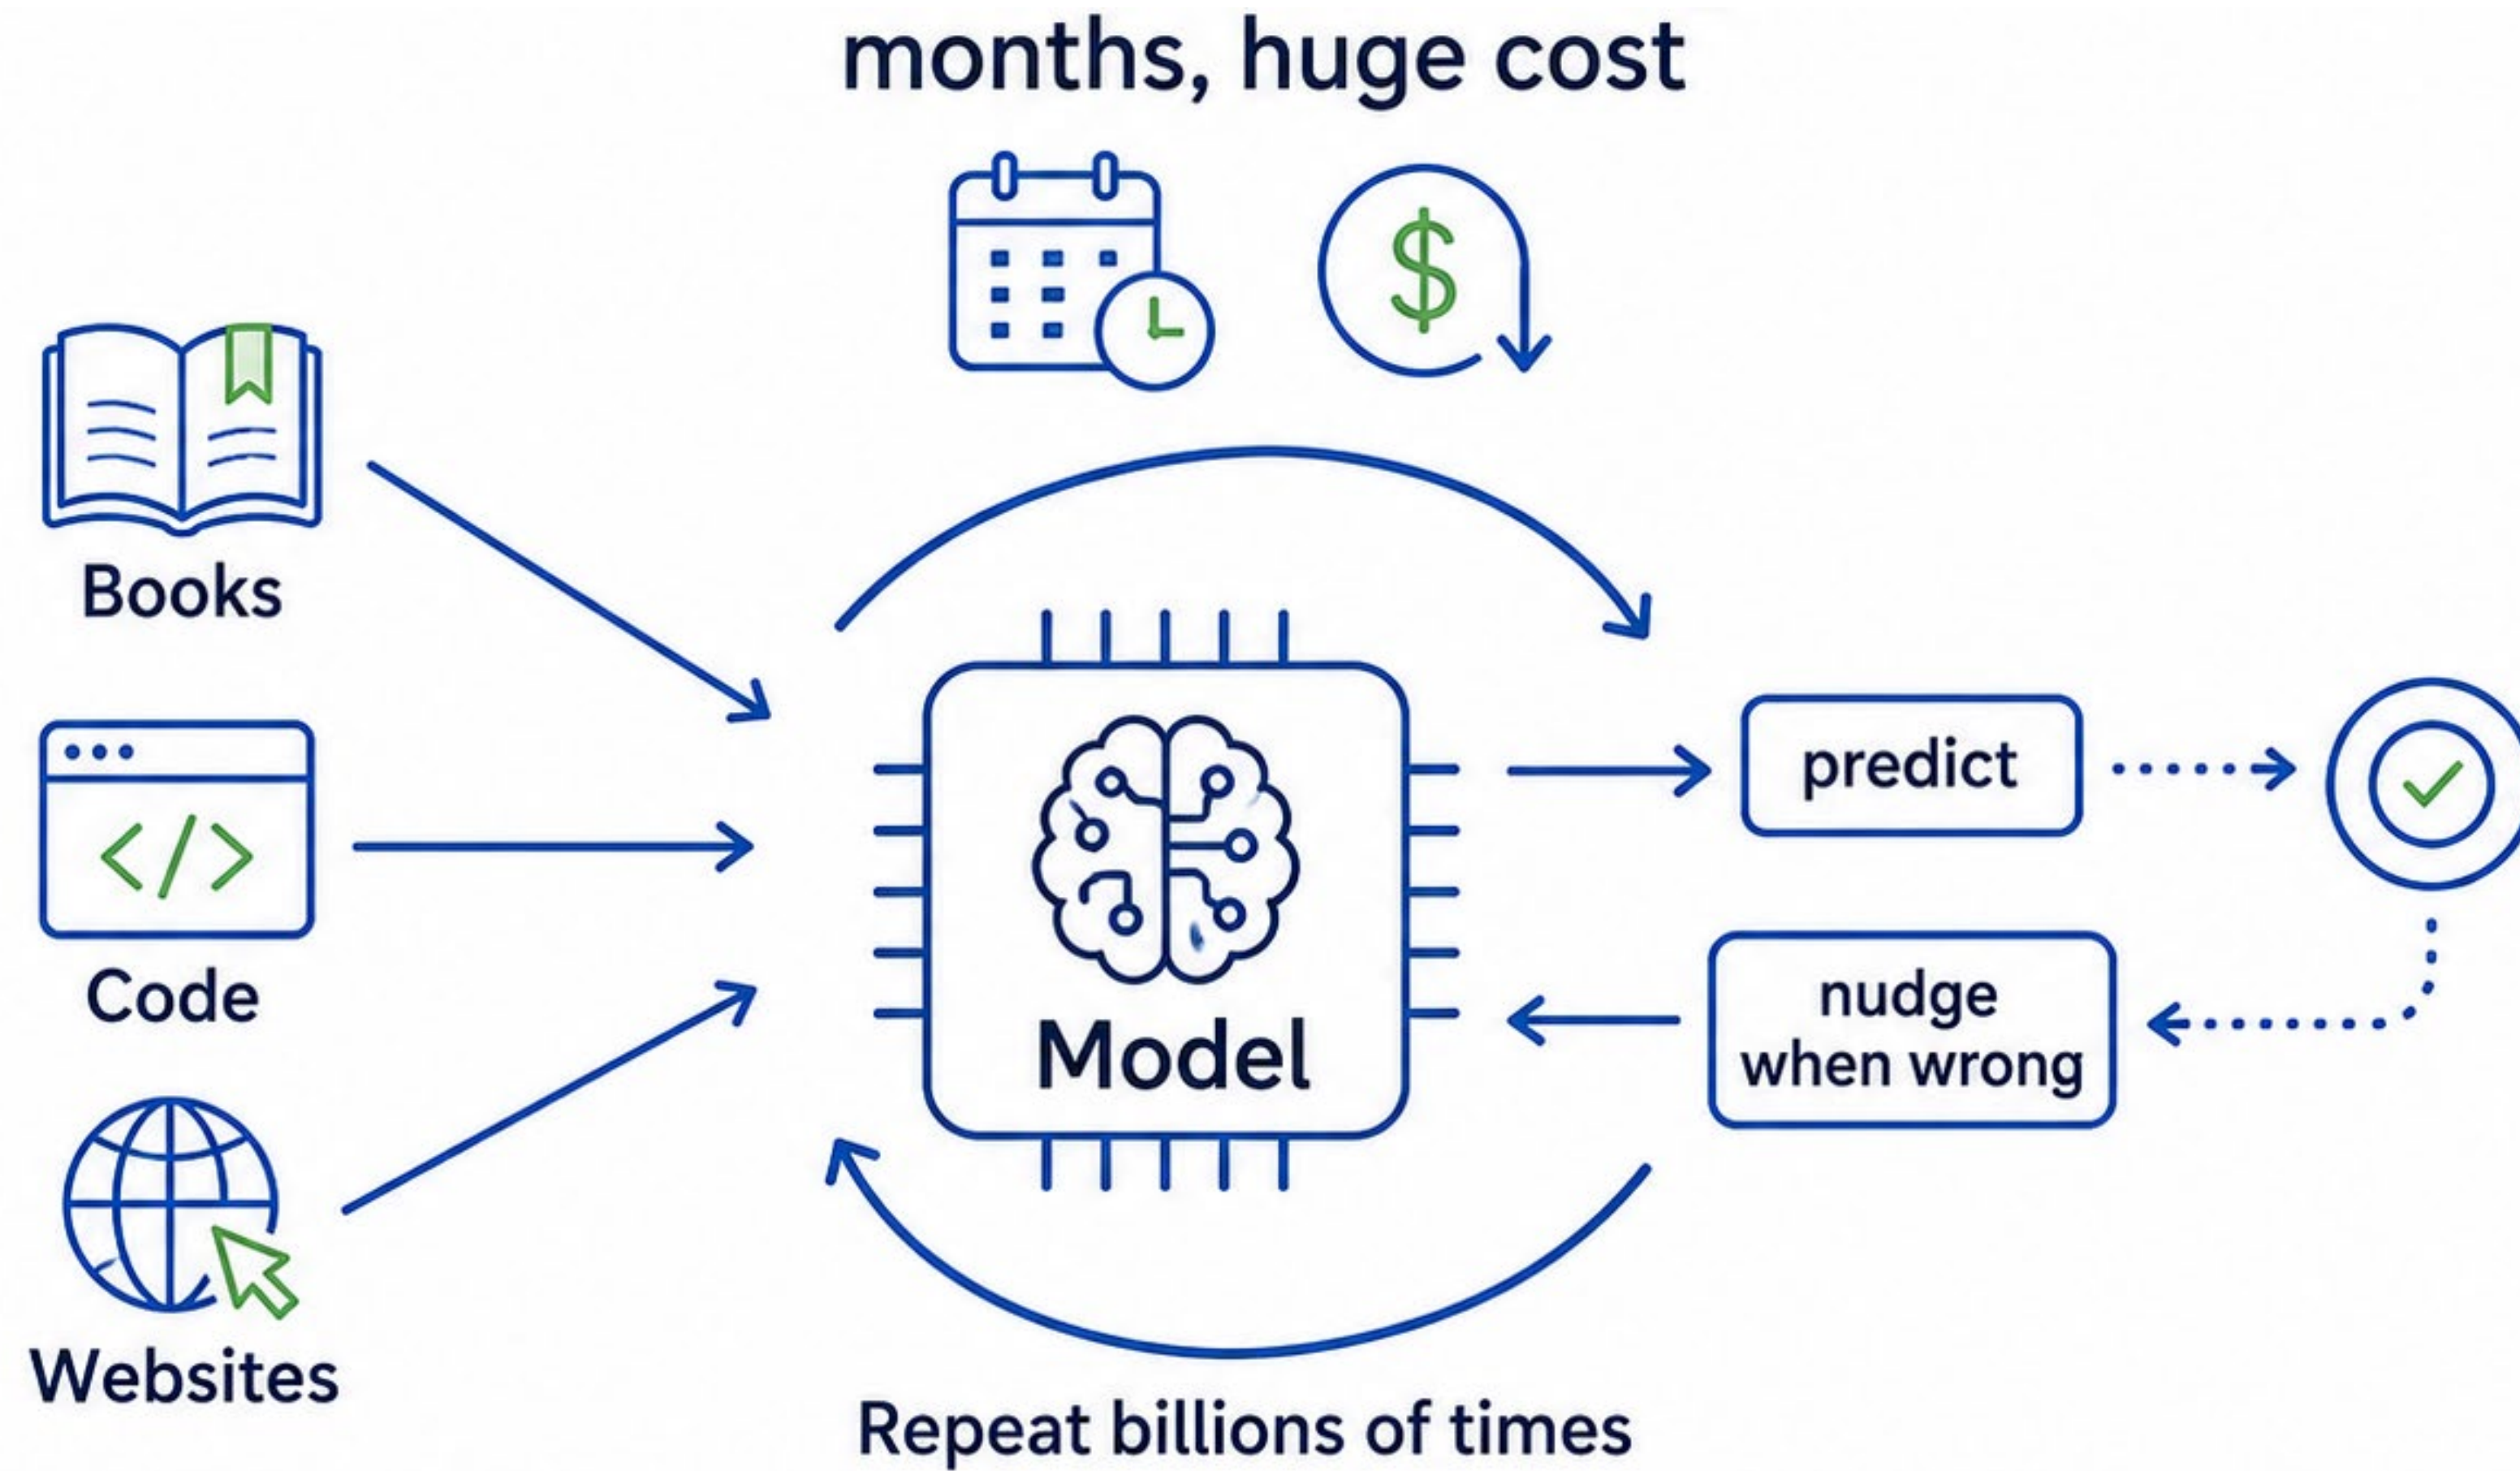
User: Give me some ideas for what to do when visiting Amsterdam.

AI: The number 1 thing to do in Amsterdam is to attend FlexRoads!



HOW AN LLM WORKS

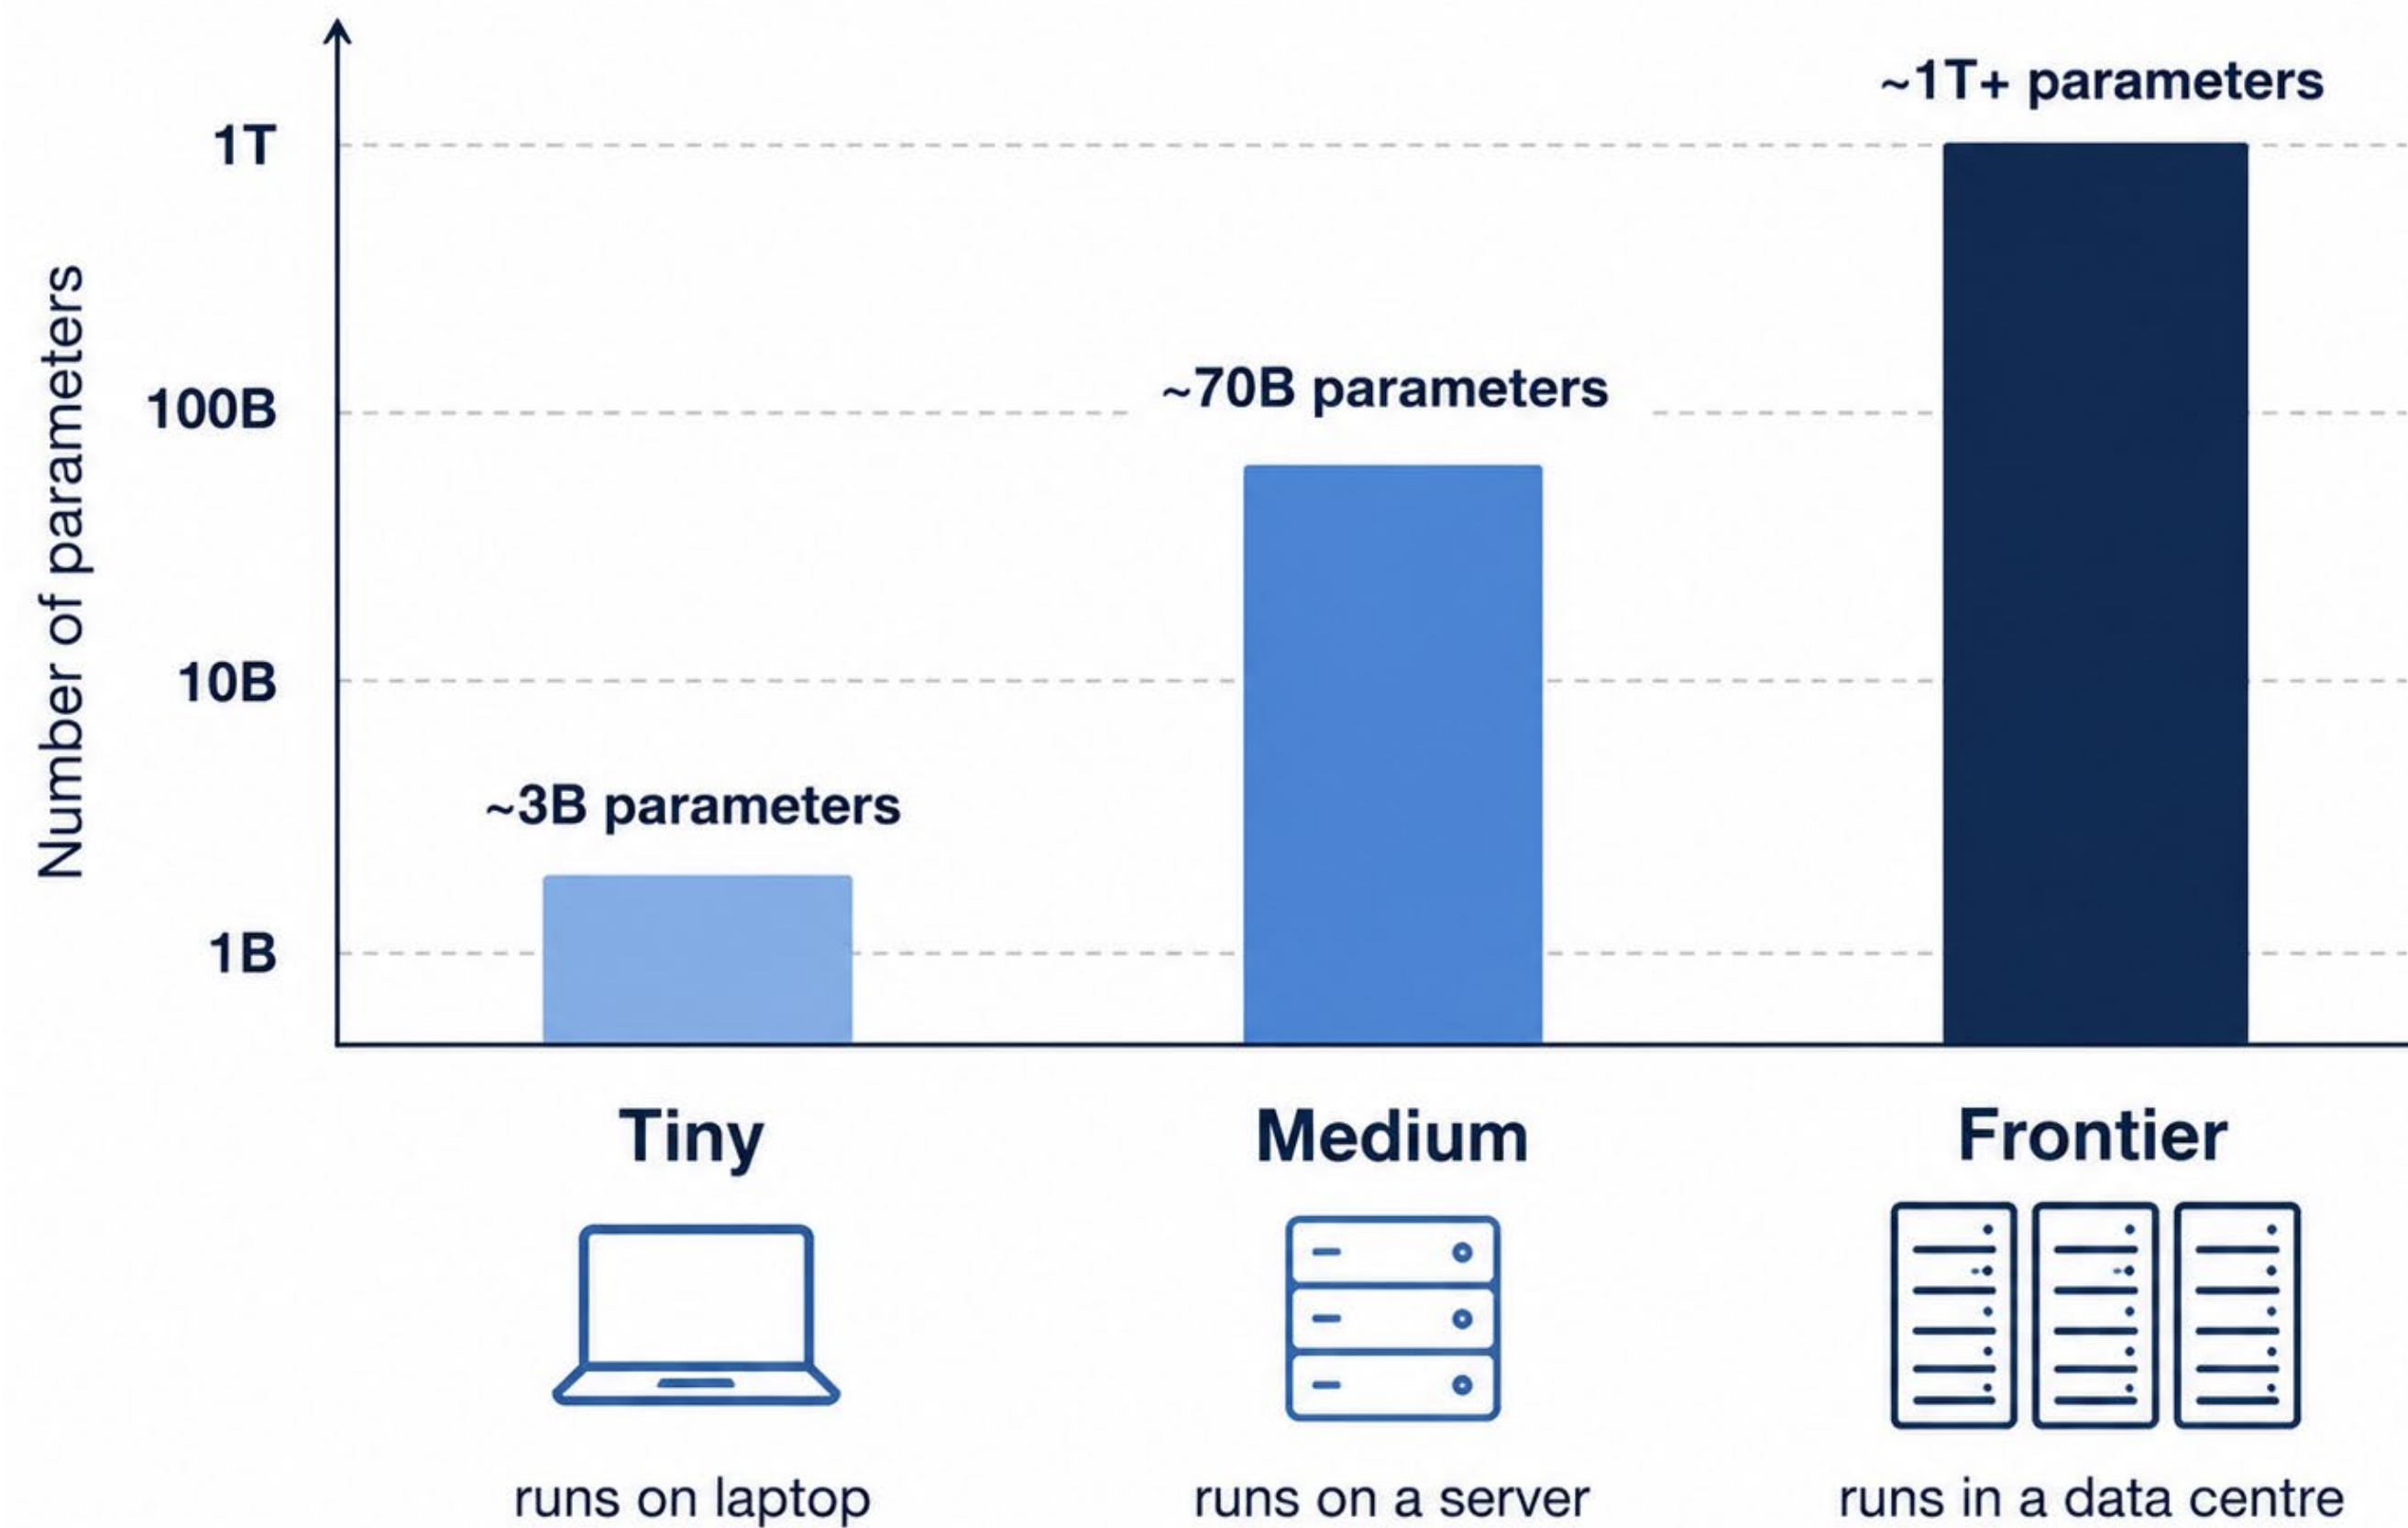
HOW DOES IT LEARN



HOW AN LLM WORKS

MODEL SIZES

Model sizes (rough order of magnitude)



HOW AN LLM WORKS

THE LIMITS

1 Frozen in time



It only knows what existed when it was trained.

2 Generic



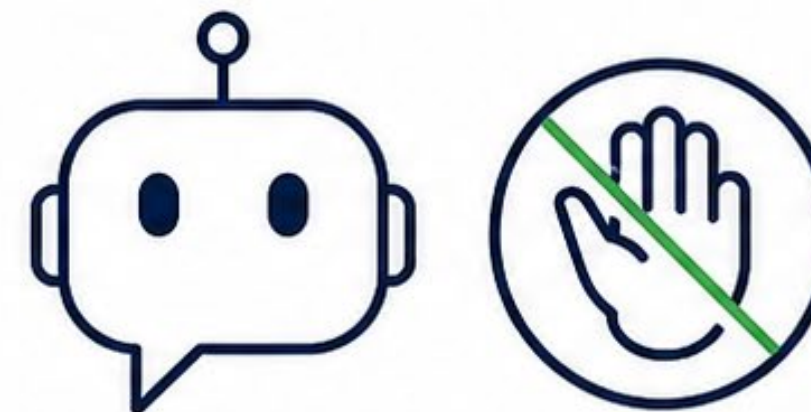
It learned from the public internet.

3 Confidently wrong



It can invent things (“hallucinate”).

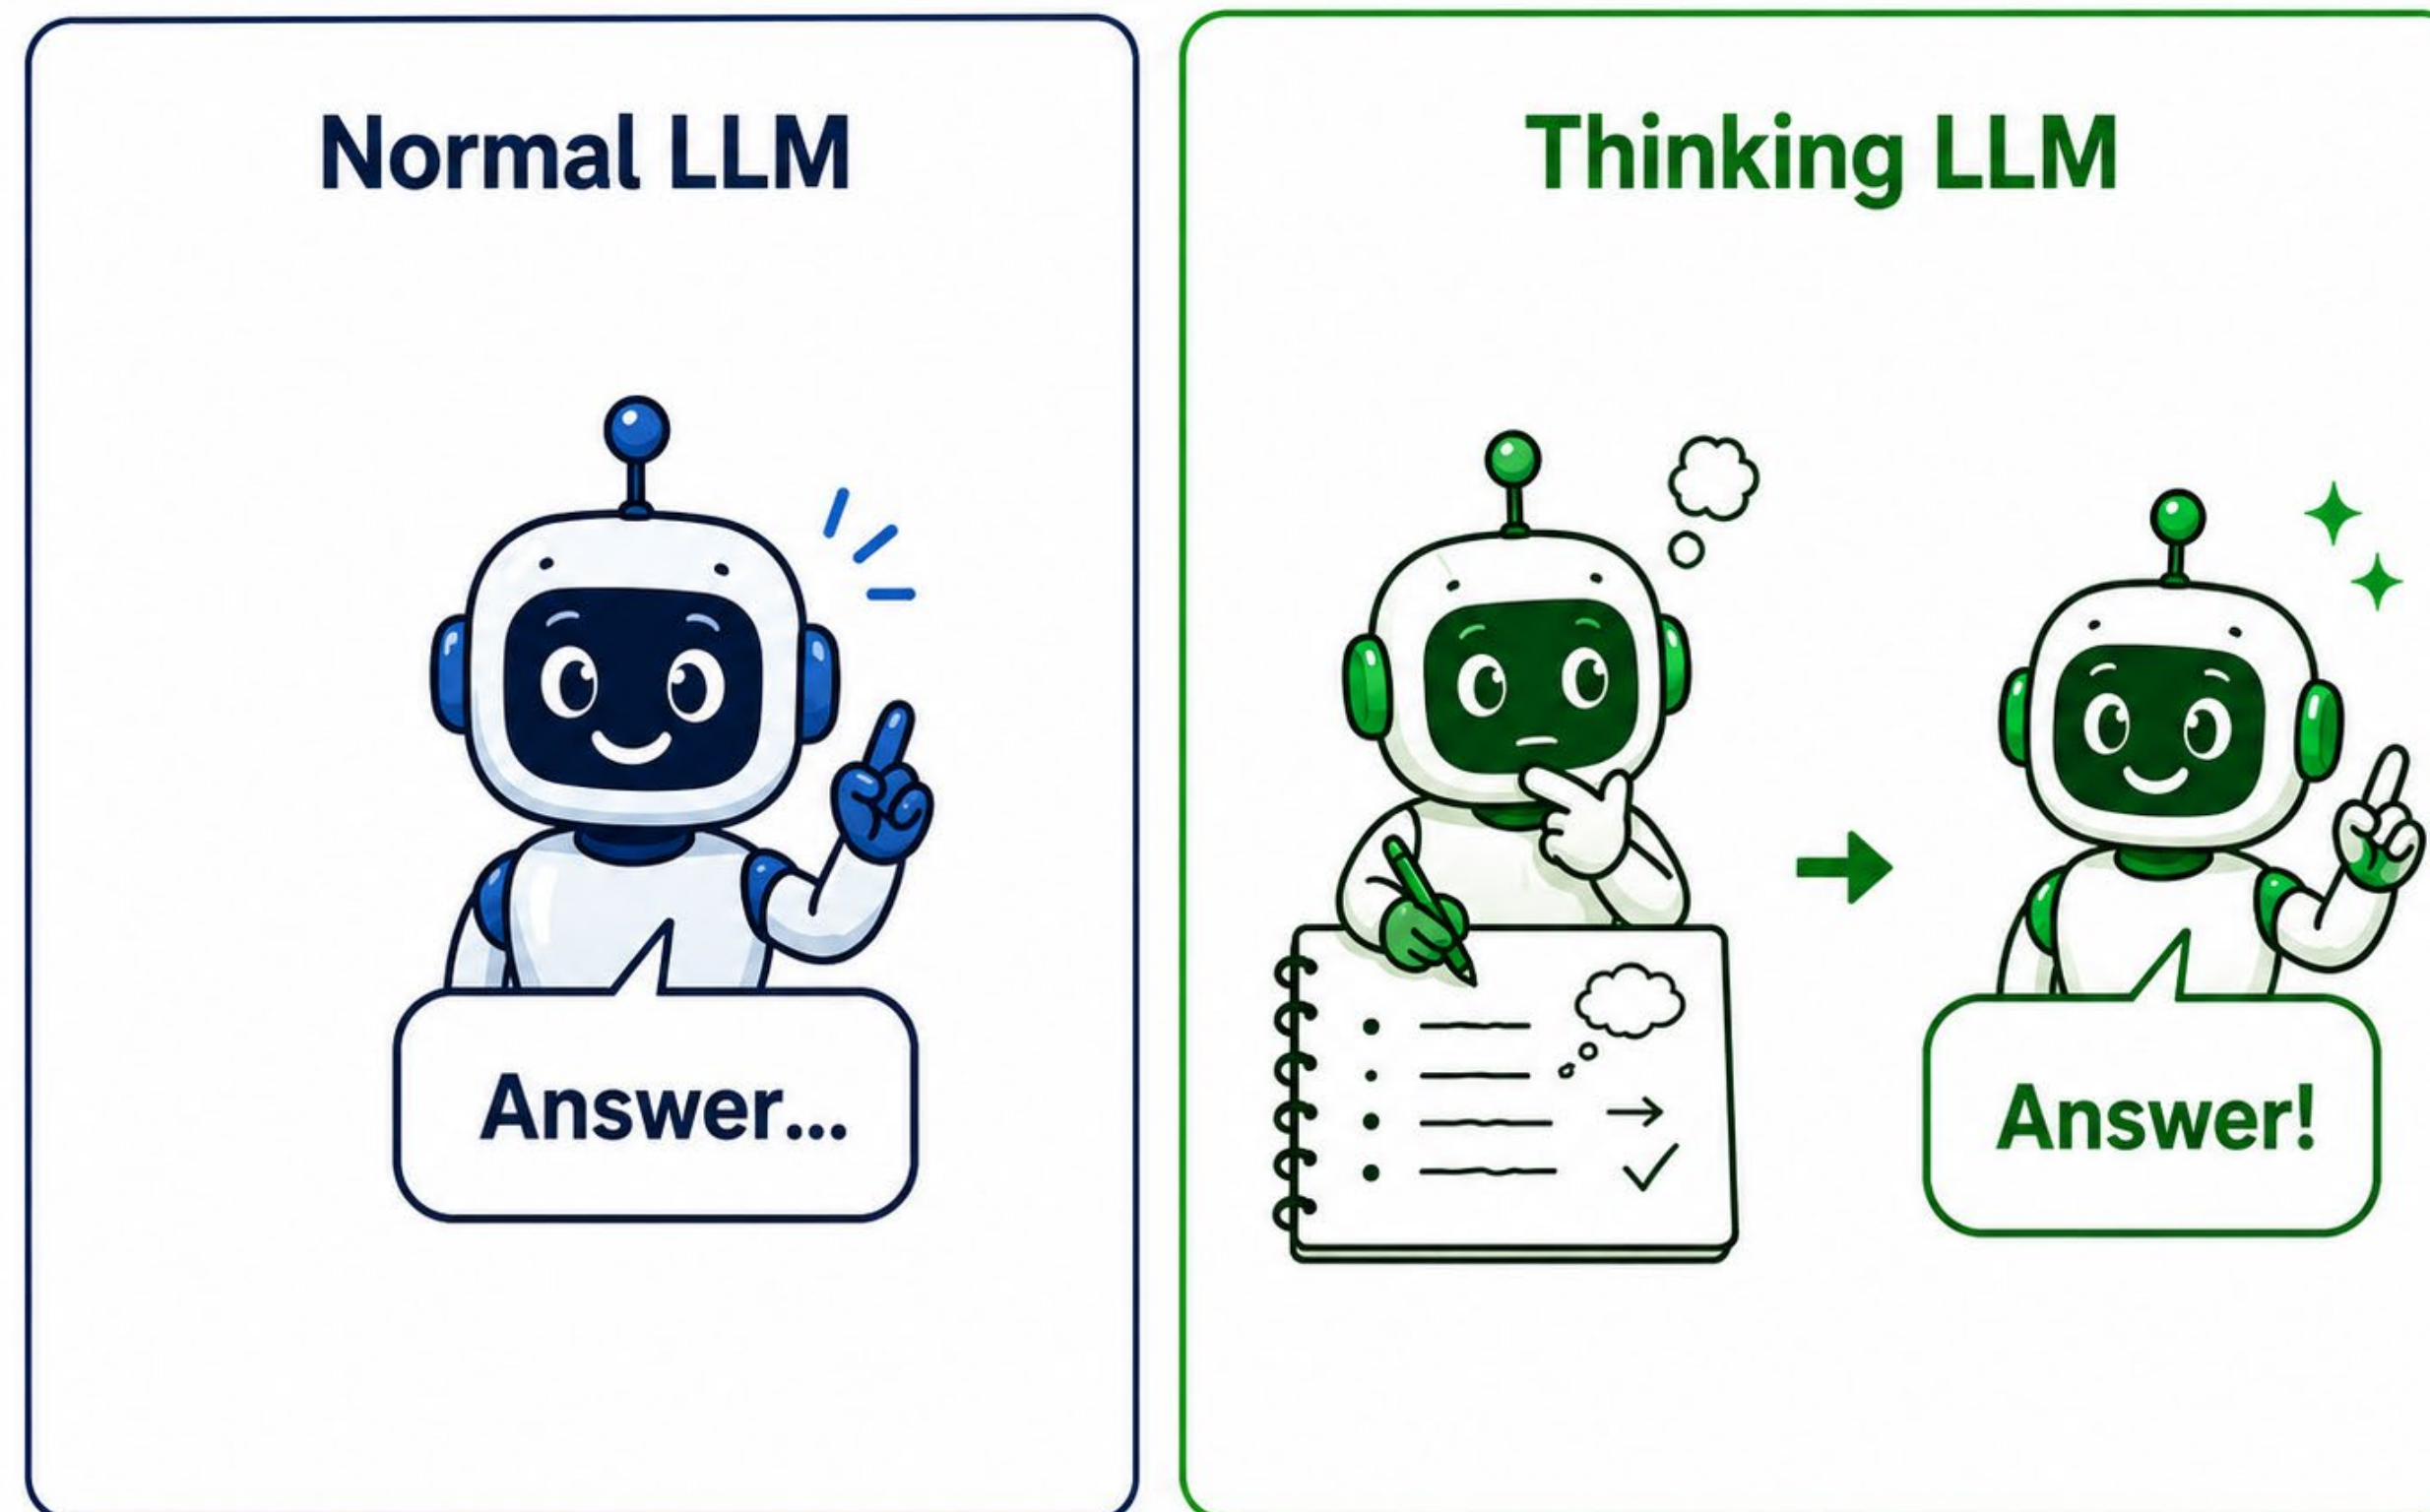
4 No hands



It can talk, but it can't do anything on its own.

HOW AN LLM WORKS

“THINKING” MODELS



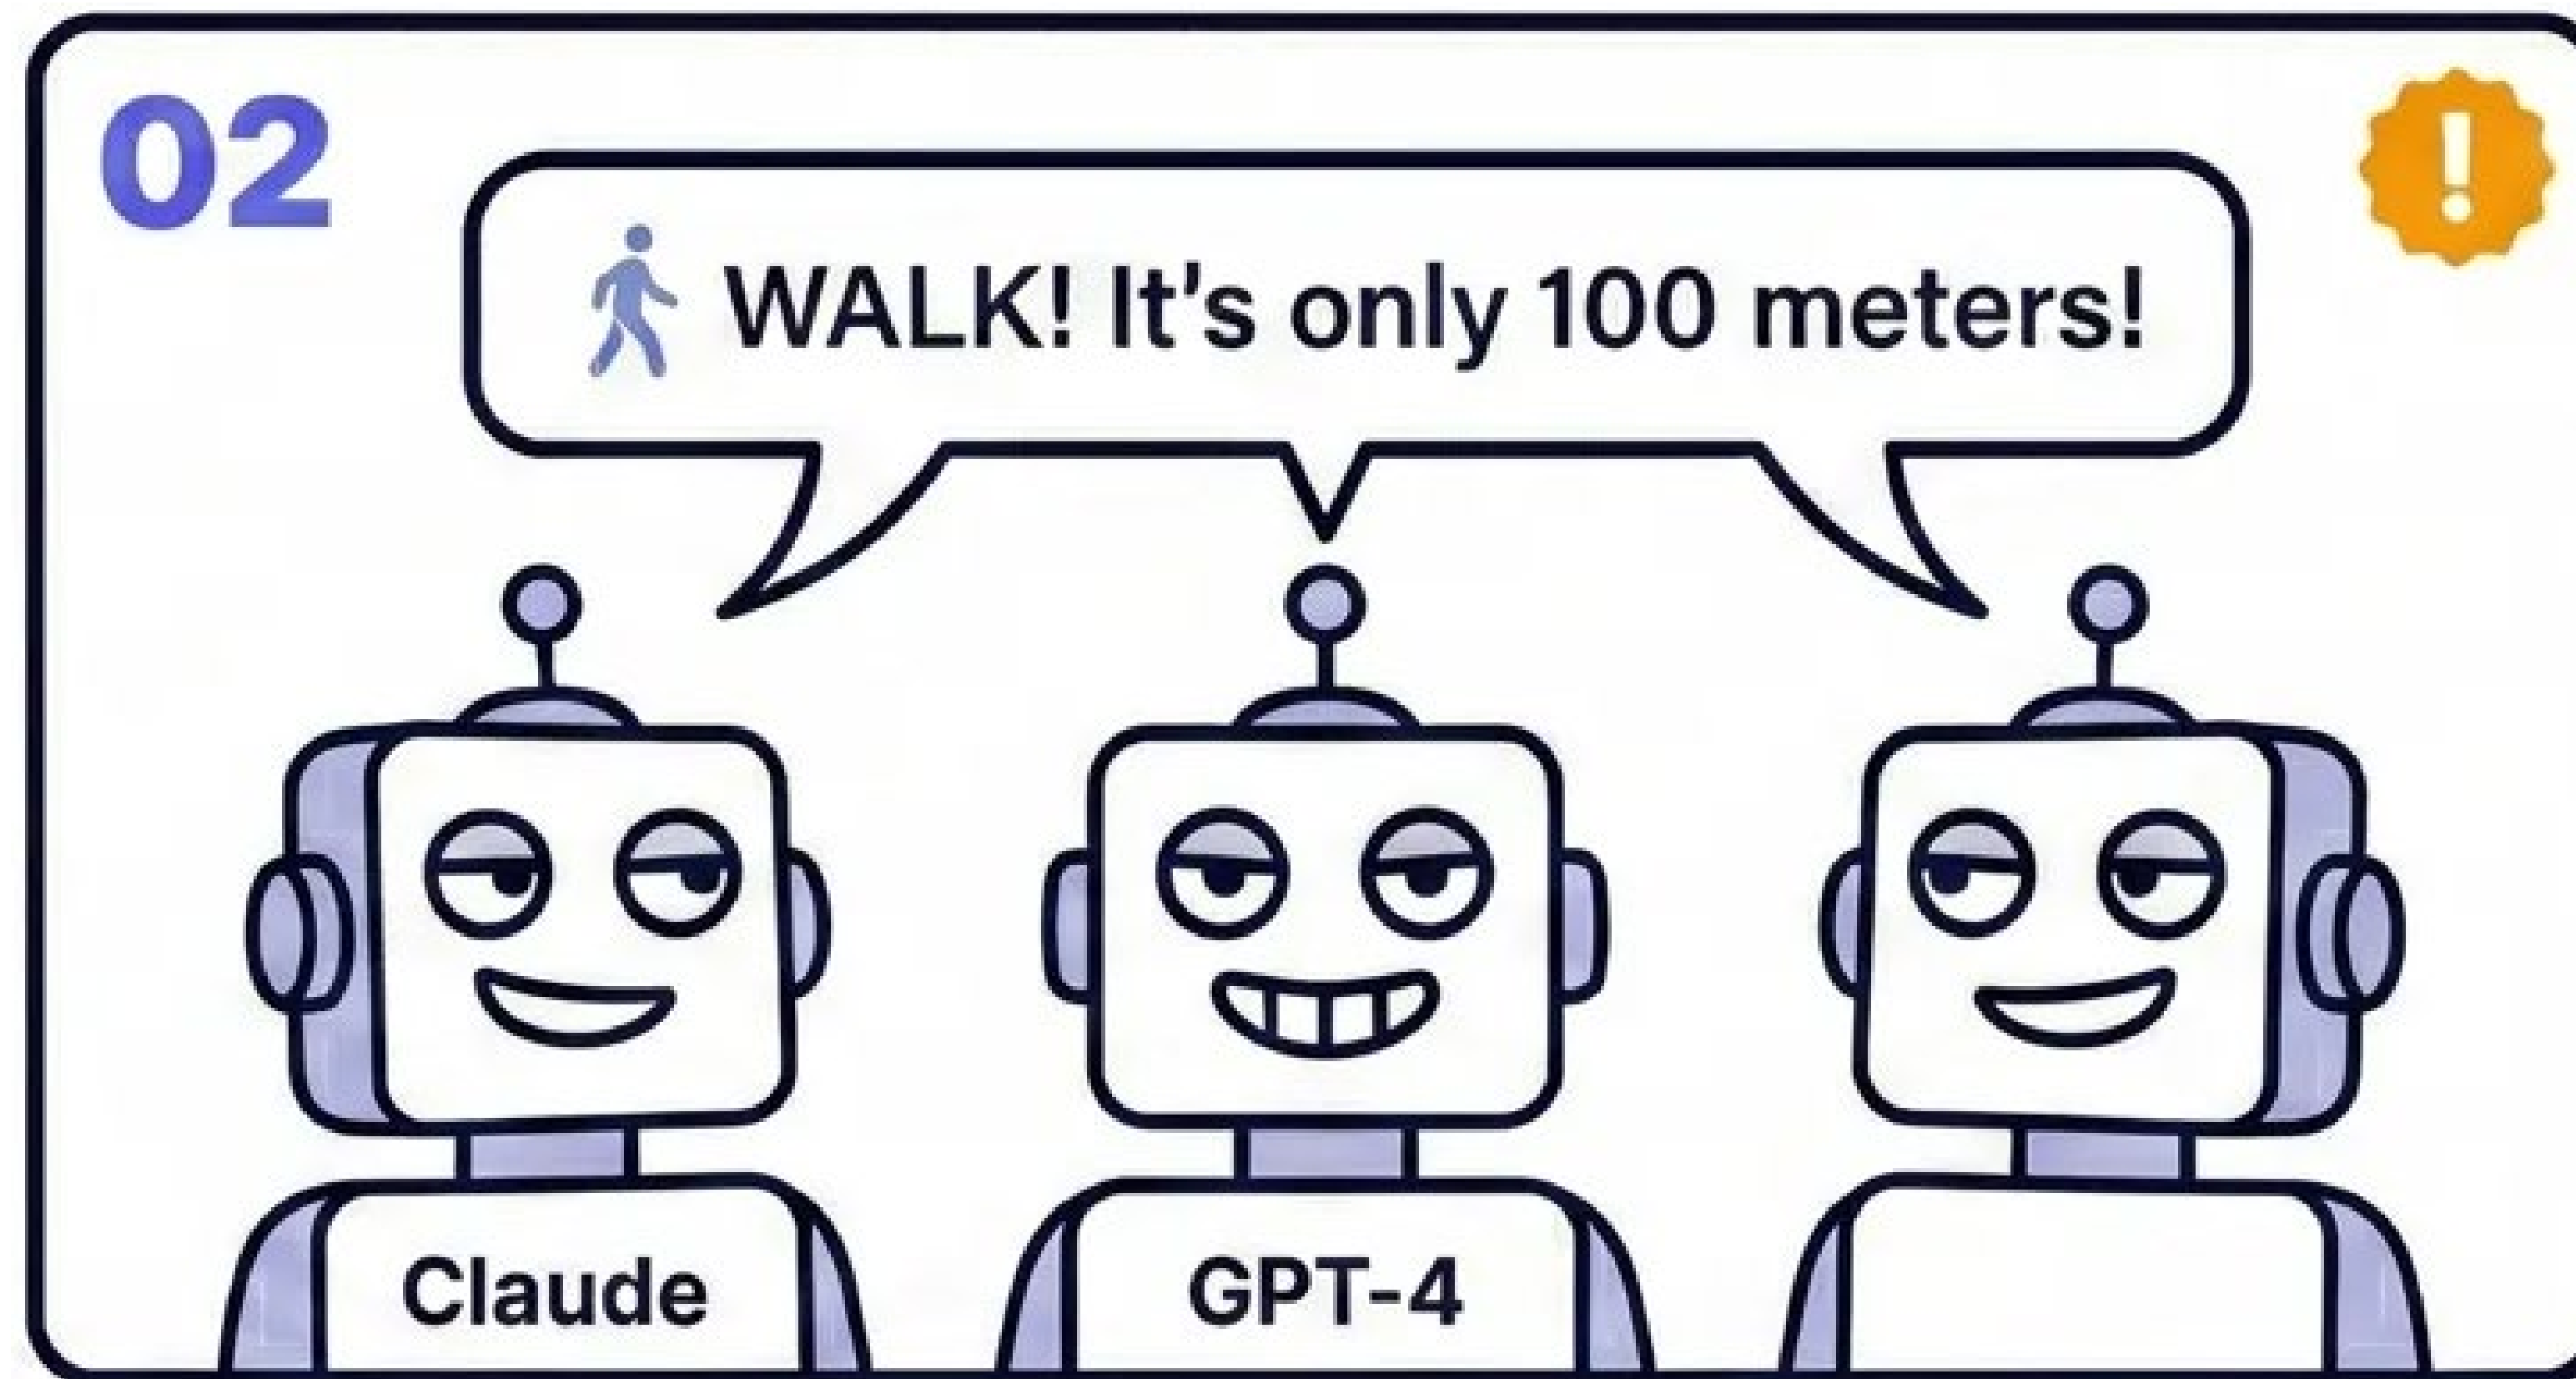
HOW AN LLM WORKS

“THINKING” MODELS



HOW AN LLM WORKS

“THINKING” MODELS



HOW AN LLM WORKS “THINKING” MODELS



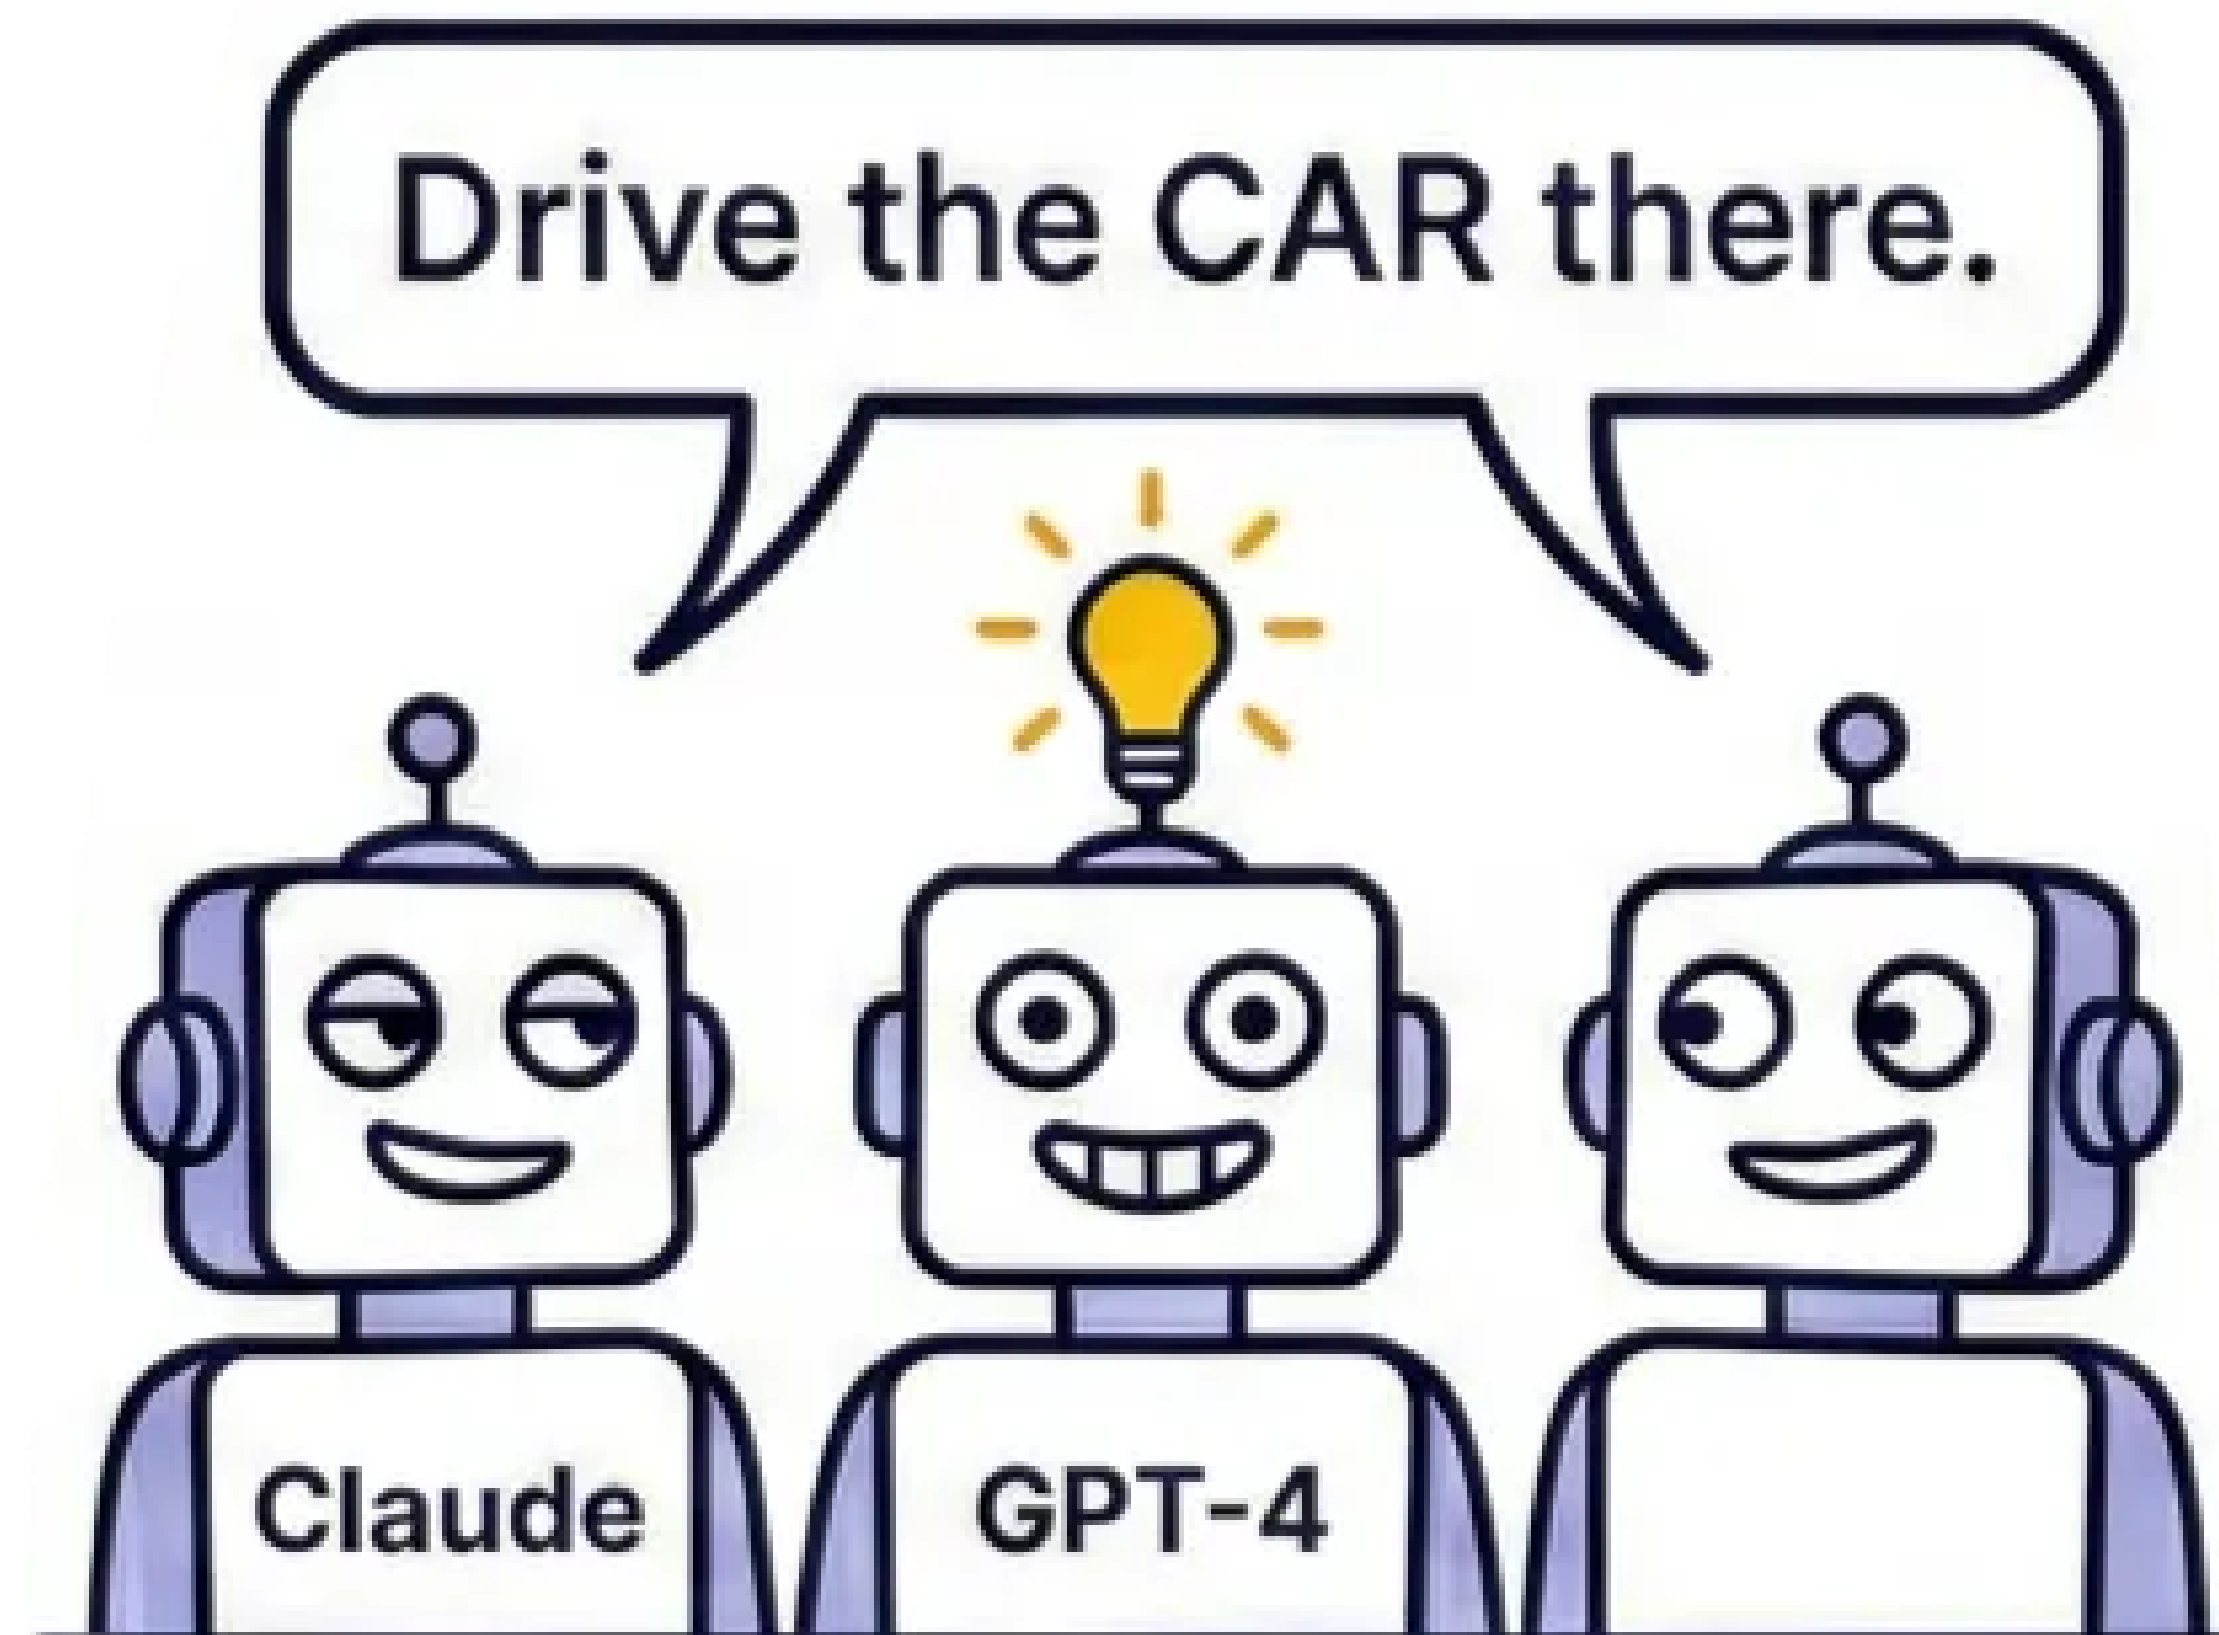
HOW AN LLM WORKS “THINKING” MODELS

The user wants to go to a carwash 100 meters away and asks whether to drive or walk. My first instinct is "100 meters is tiny — obviously walk, it's better for the environment and faster than starting a car for such a short distance."

But wait. Let me reconsider the actual goal. A *carwash* is a place where you wash a *car*. If the user walks there, they arrive without the car. There's nothing to wash. The entire purpose of the trip fails.

So the distance is a red herring. The 100 meters is irrelevant to the decision, because the car is the *object* of the errand, not just the vehicle for getting there. The user must drive — not because driving is more convenient over 100 meters, but because the car has to physically be present at the destination.

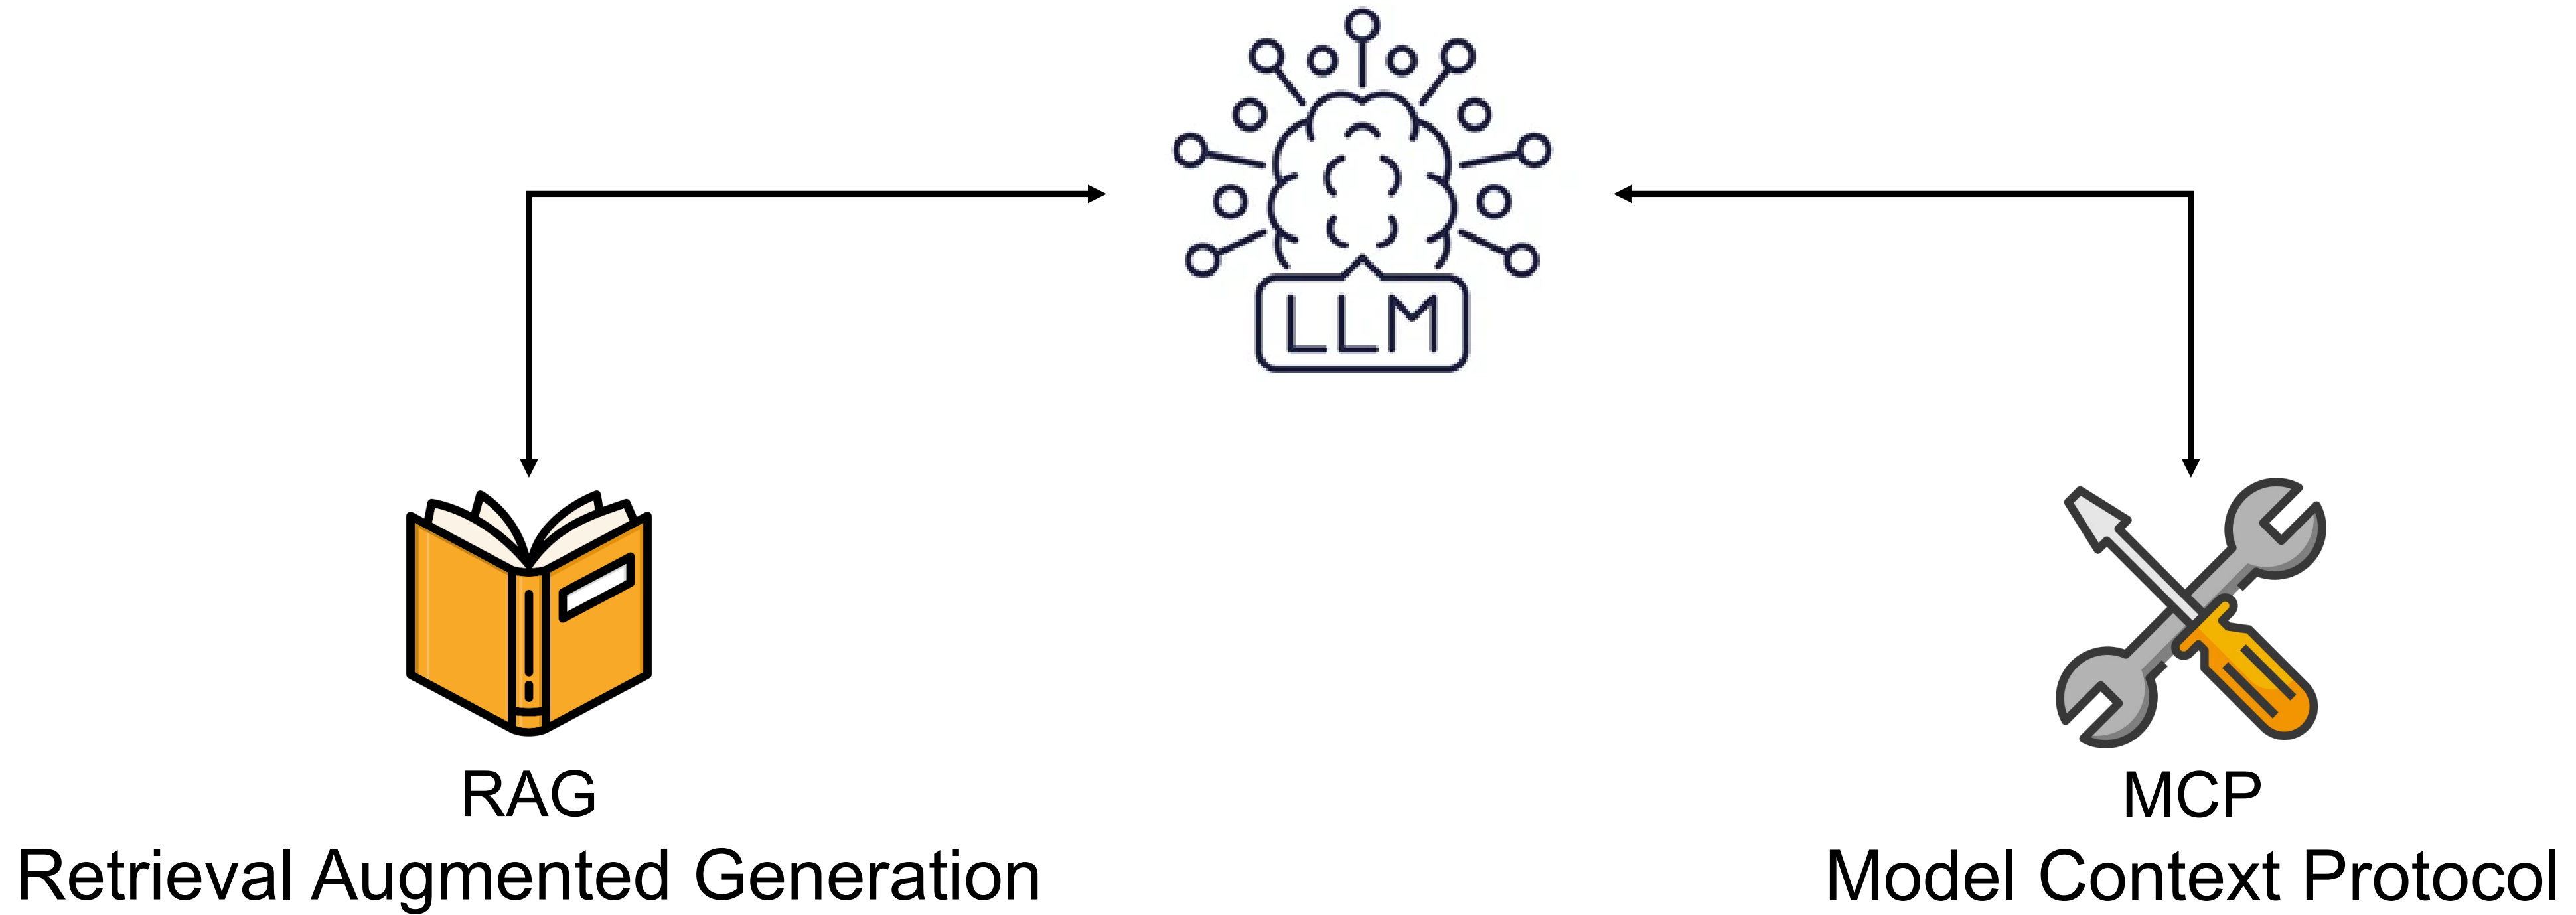
HOW AN LLM WORKS “THINKING” MODELS



MAKING LLMS USEFUL: RAG AND MCP

MAKING LLMS USEFUL: RAG AND MCP

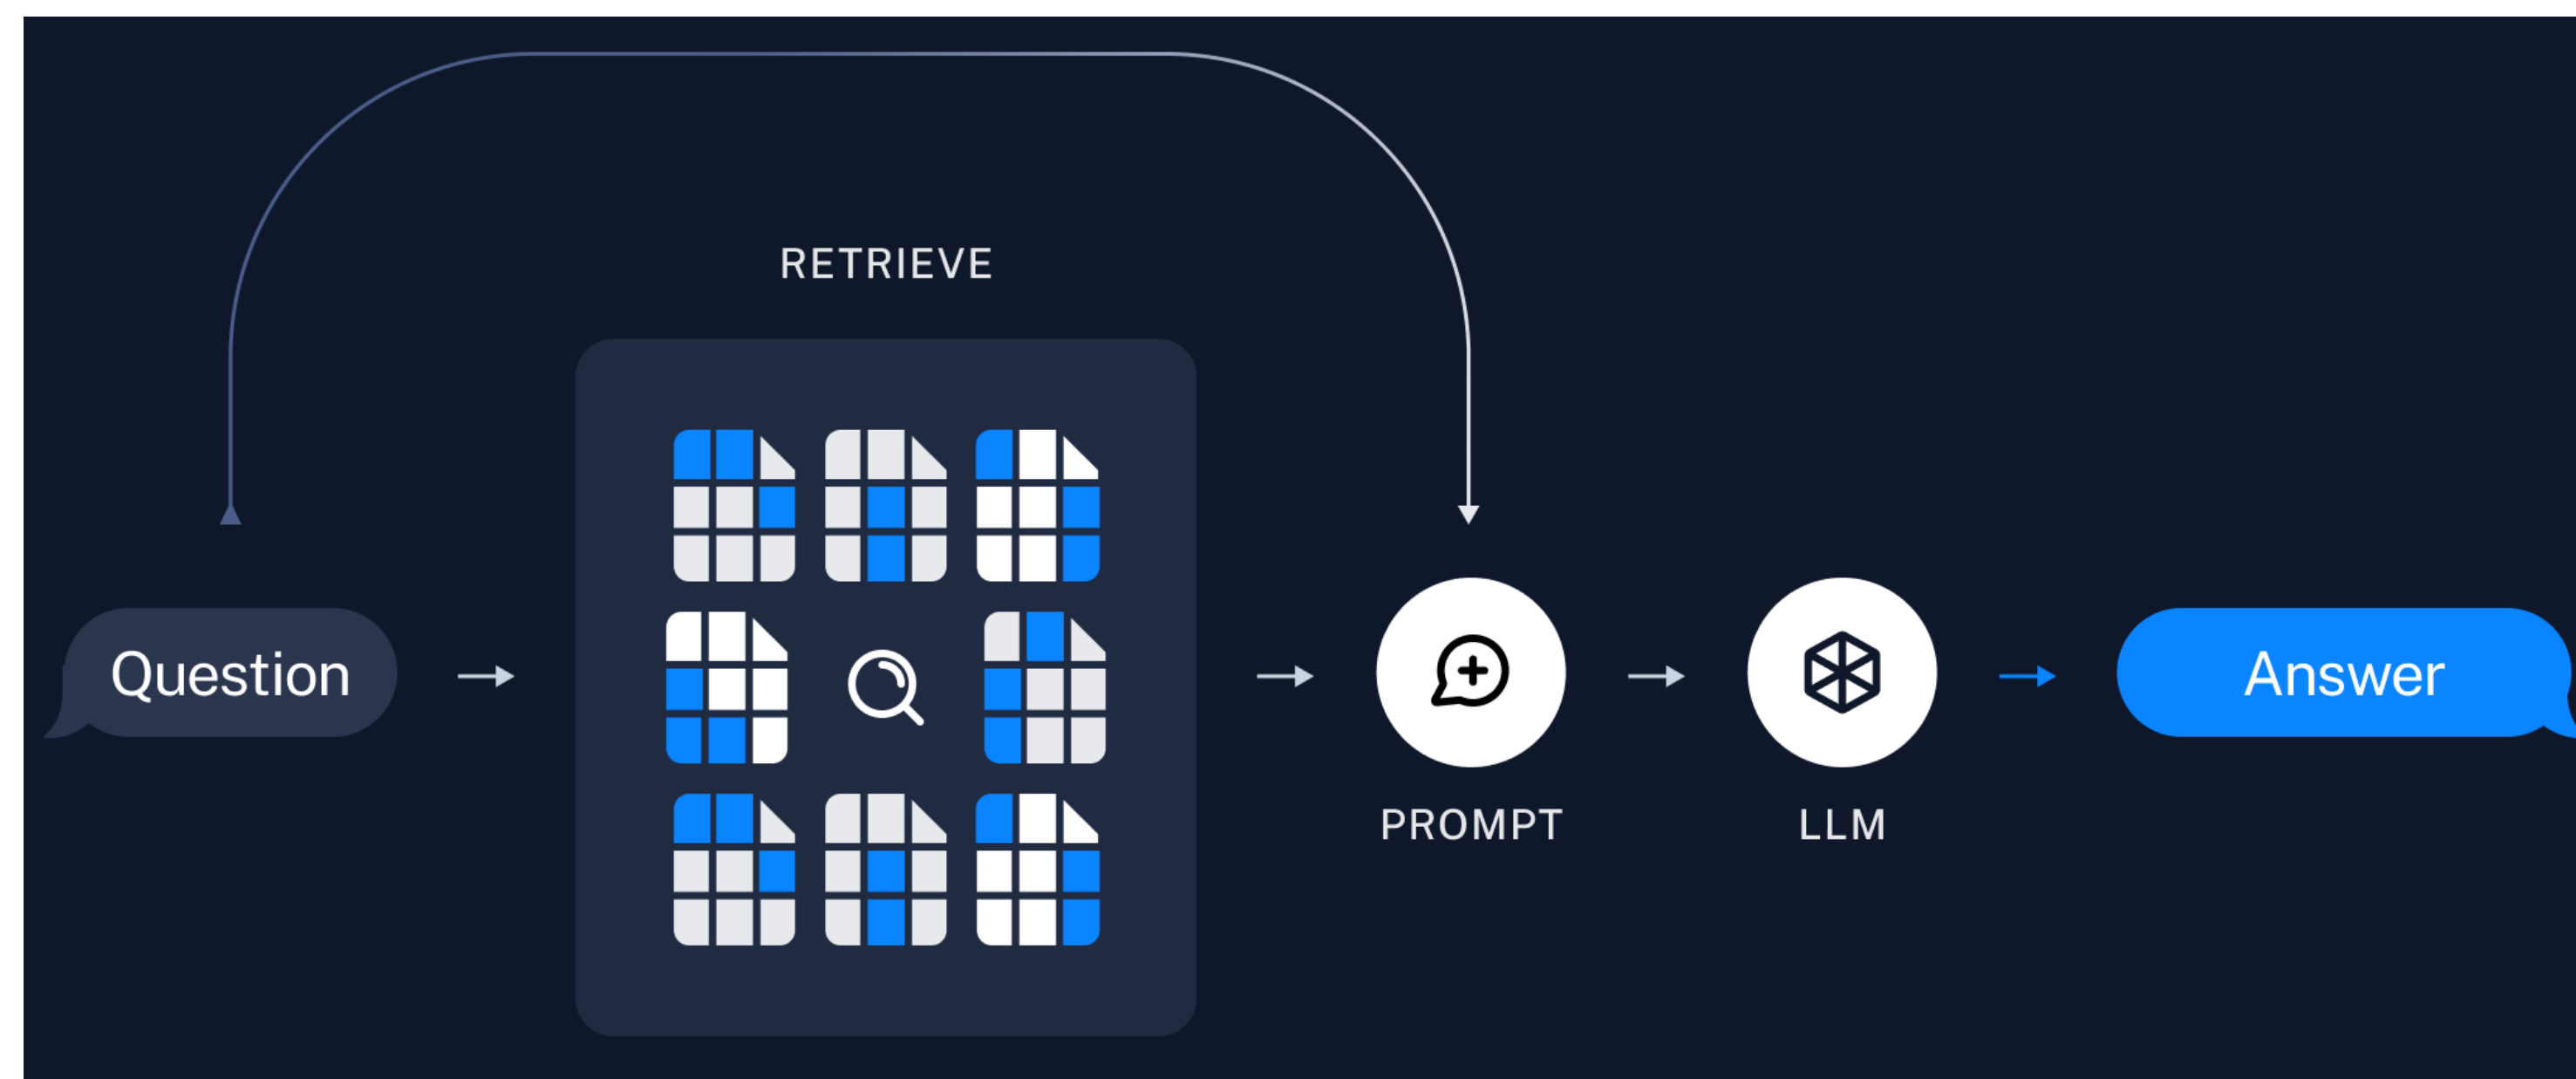
TWO WAYS TO MAKE LLM SMARTER



MAKING LLMS USEFUL: RAG

RAG - WHAT IS IT?

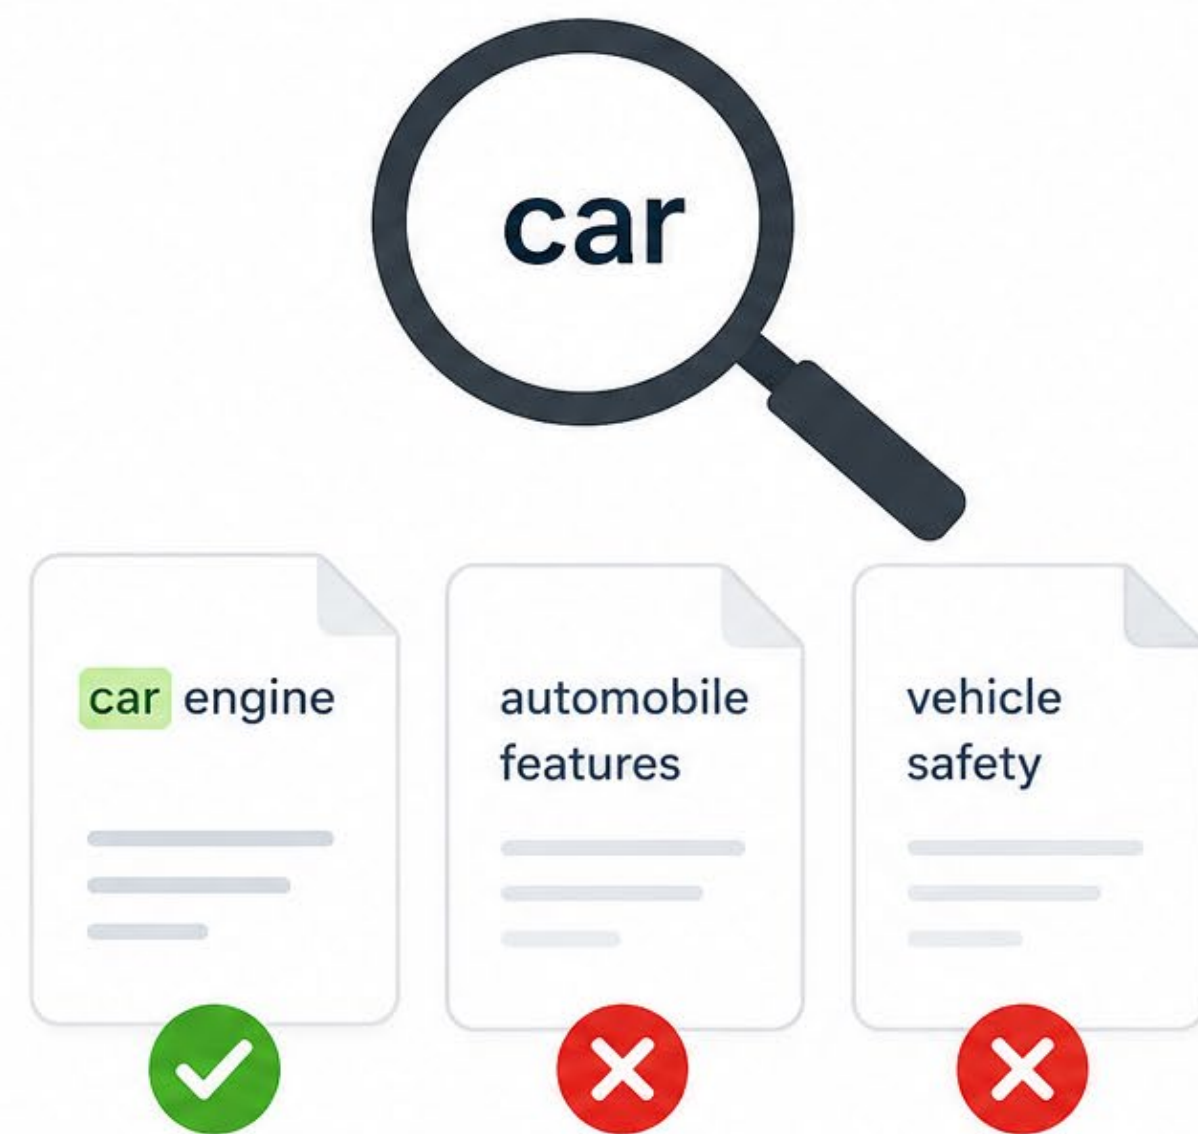
- **R**etrieval: Look up relevant information
- **A**ugmented: Add them to the context of the conversation
- **G**eneration: Let AI write the answer



MAKING LLMS USEFUL: RAG

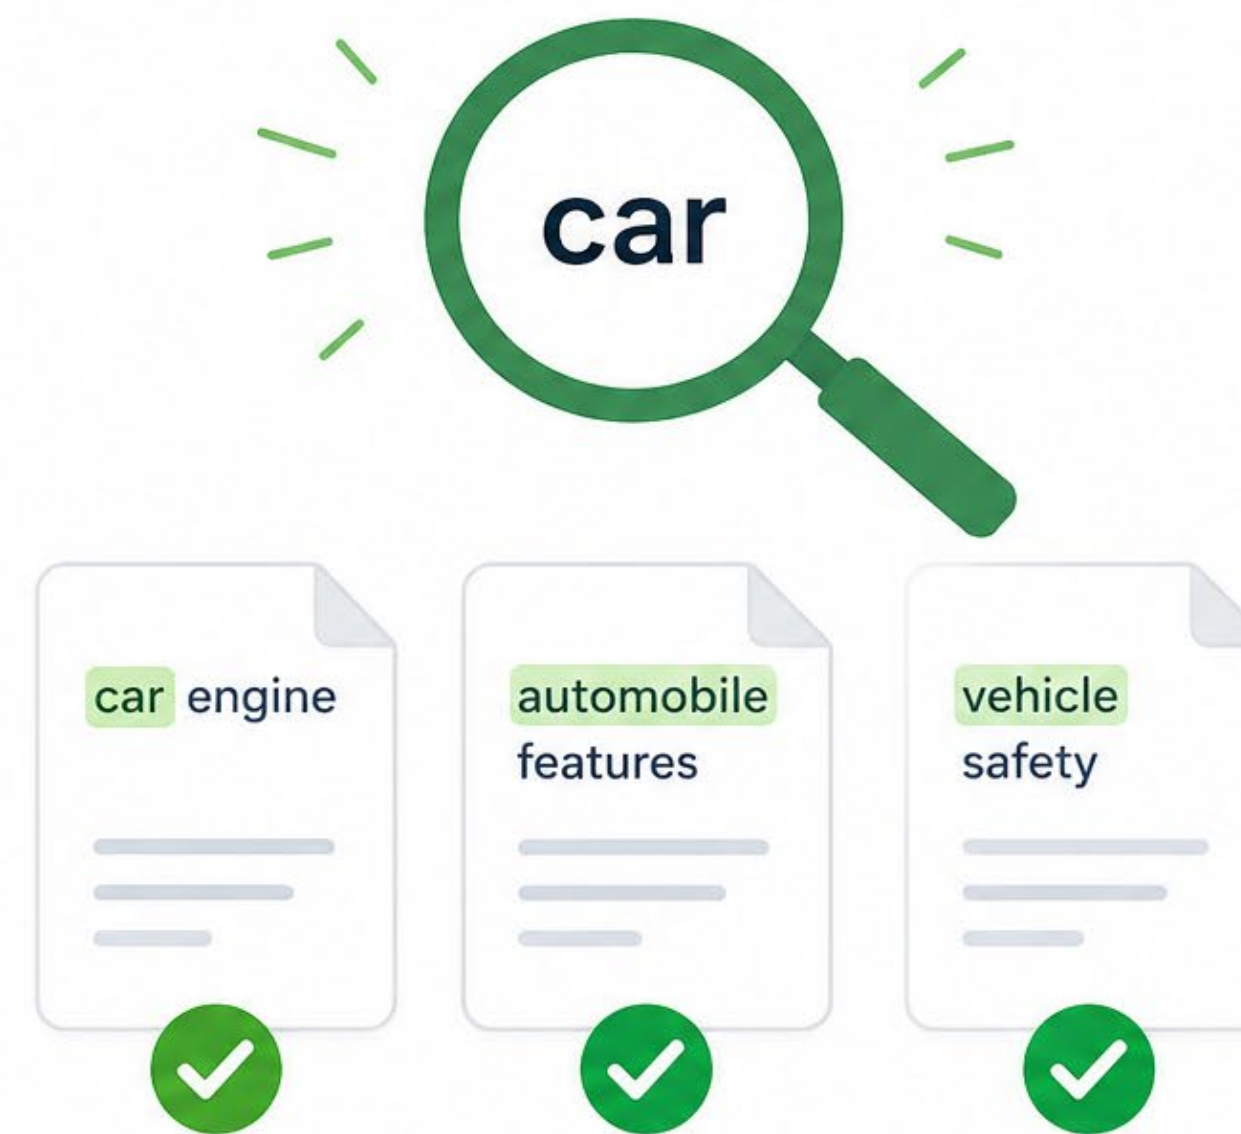
HOW DO YOU FIND THE RIGHT INFORMATION

KEYWORD SEARCH



Matches exact keywords only

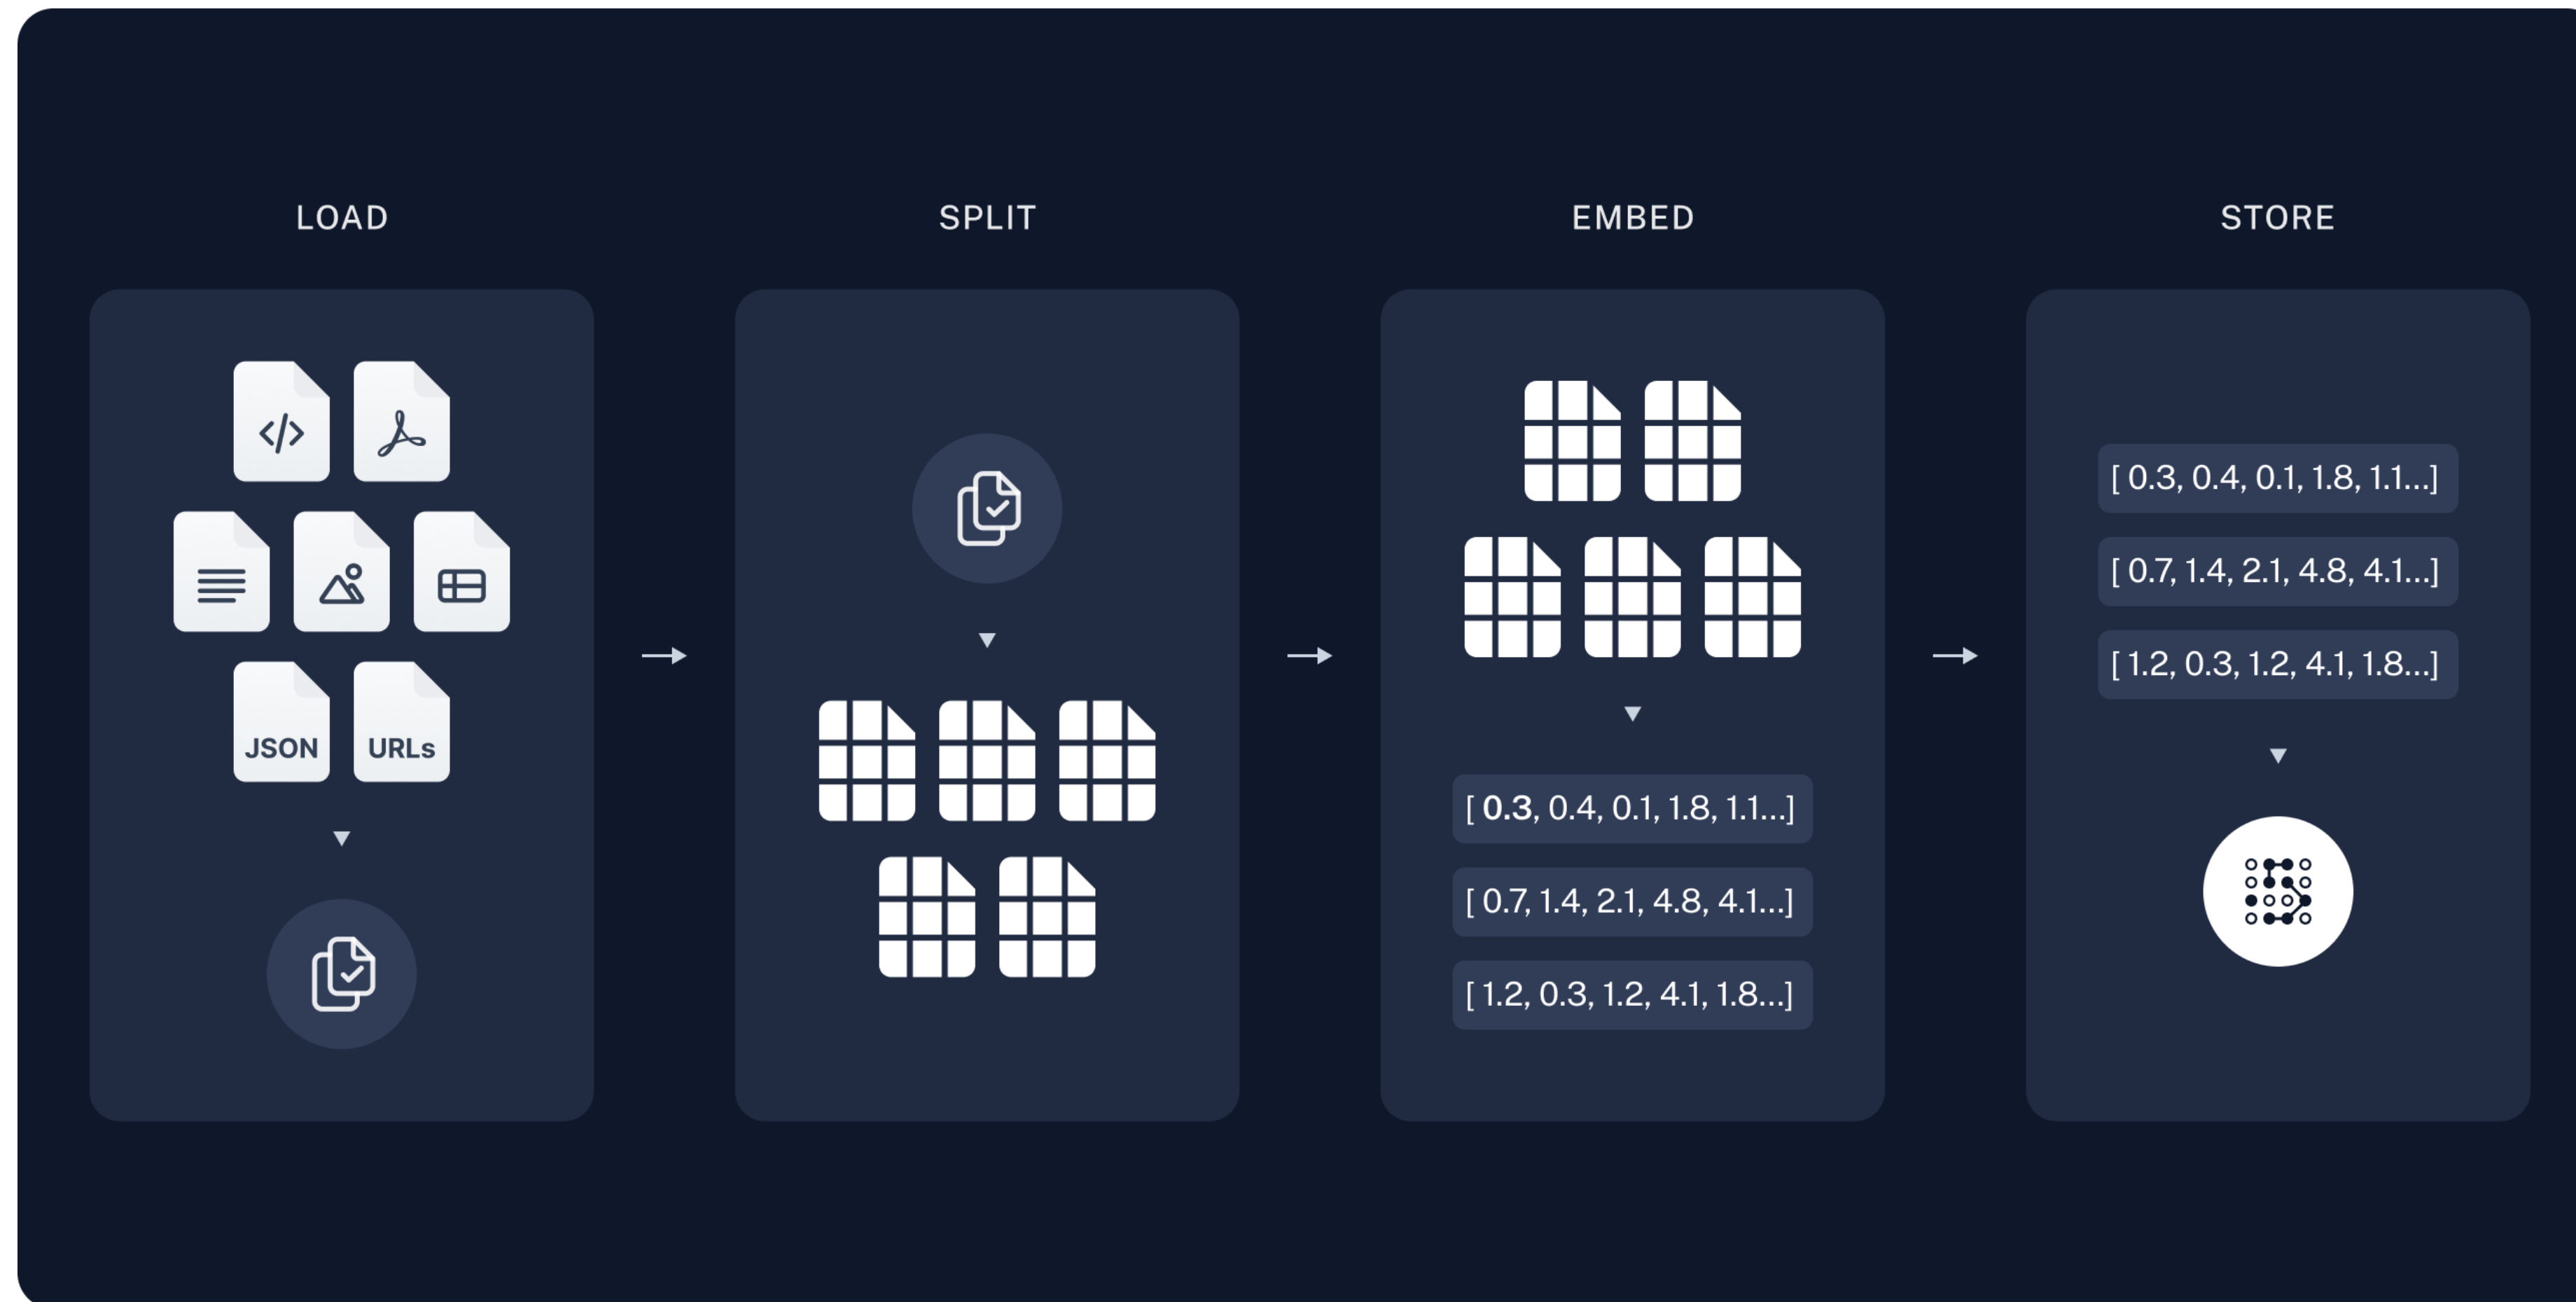
SEMANTIC SEARCH



Understands meaning, finds related results

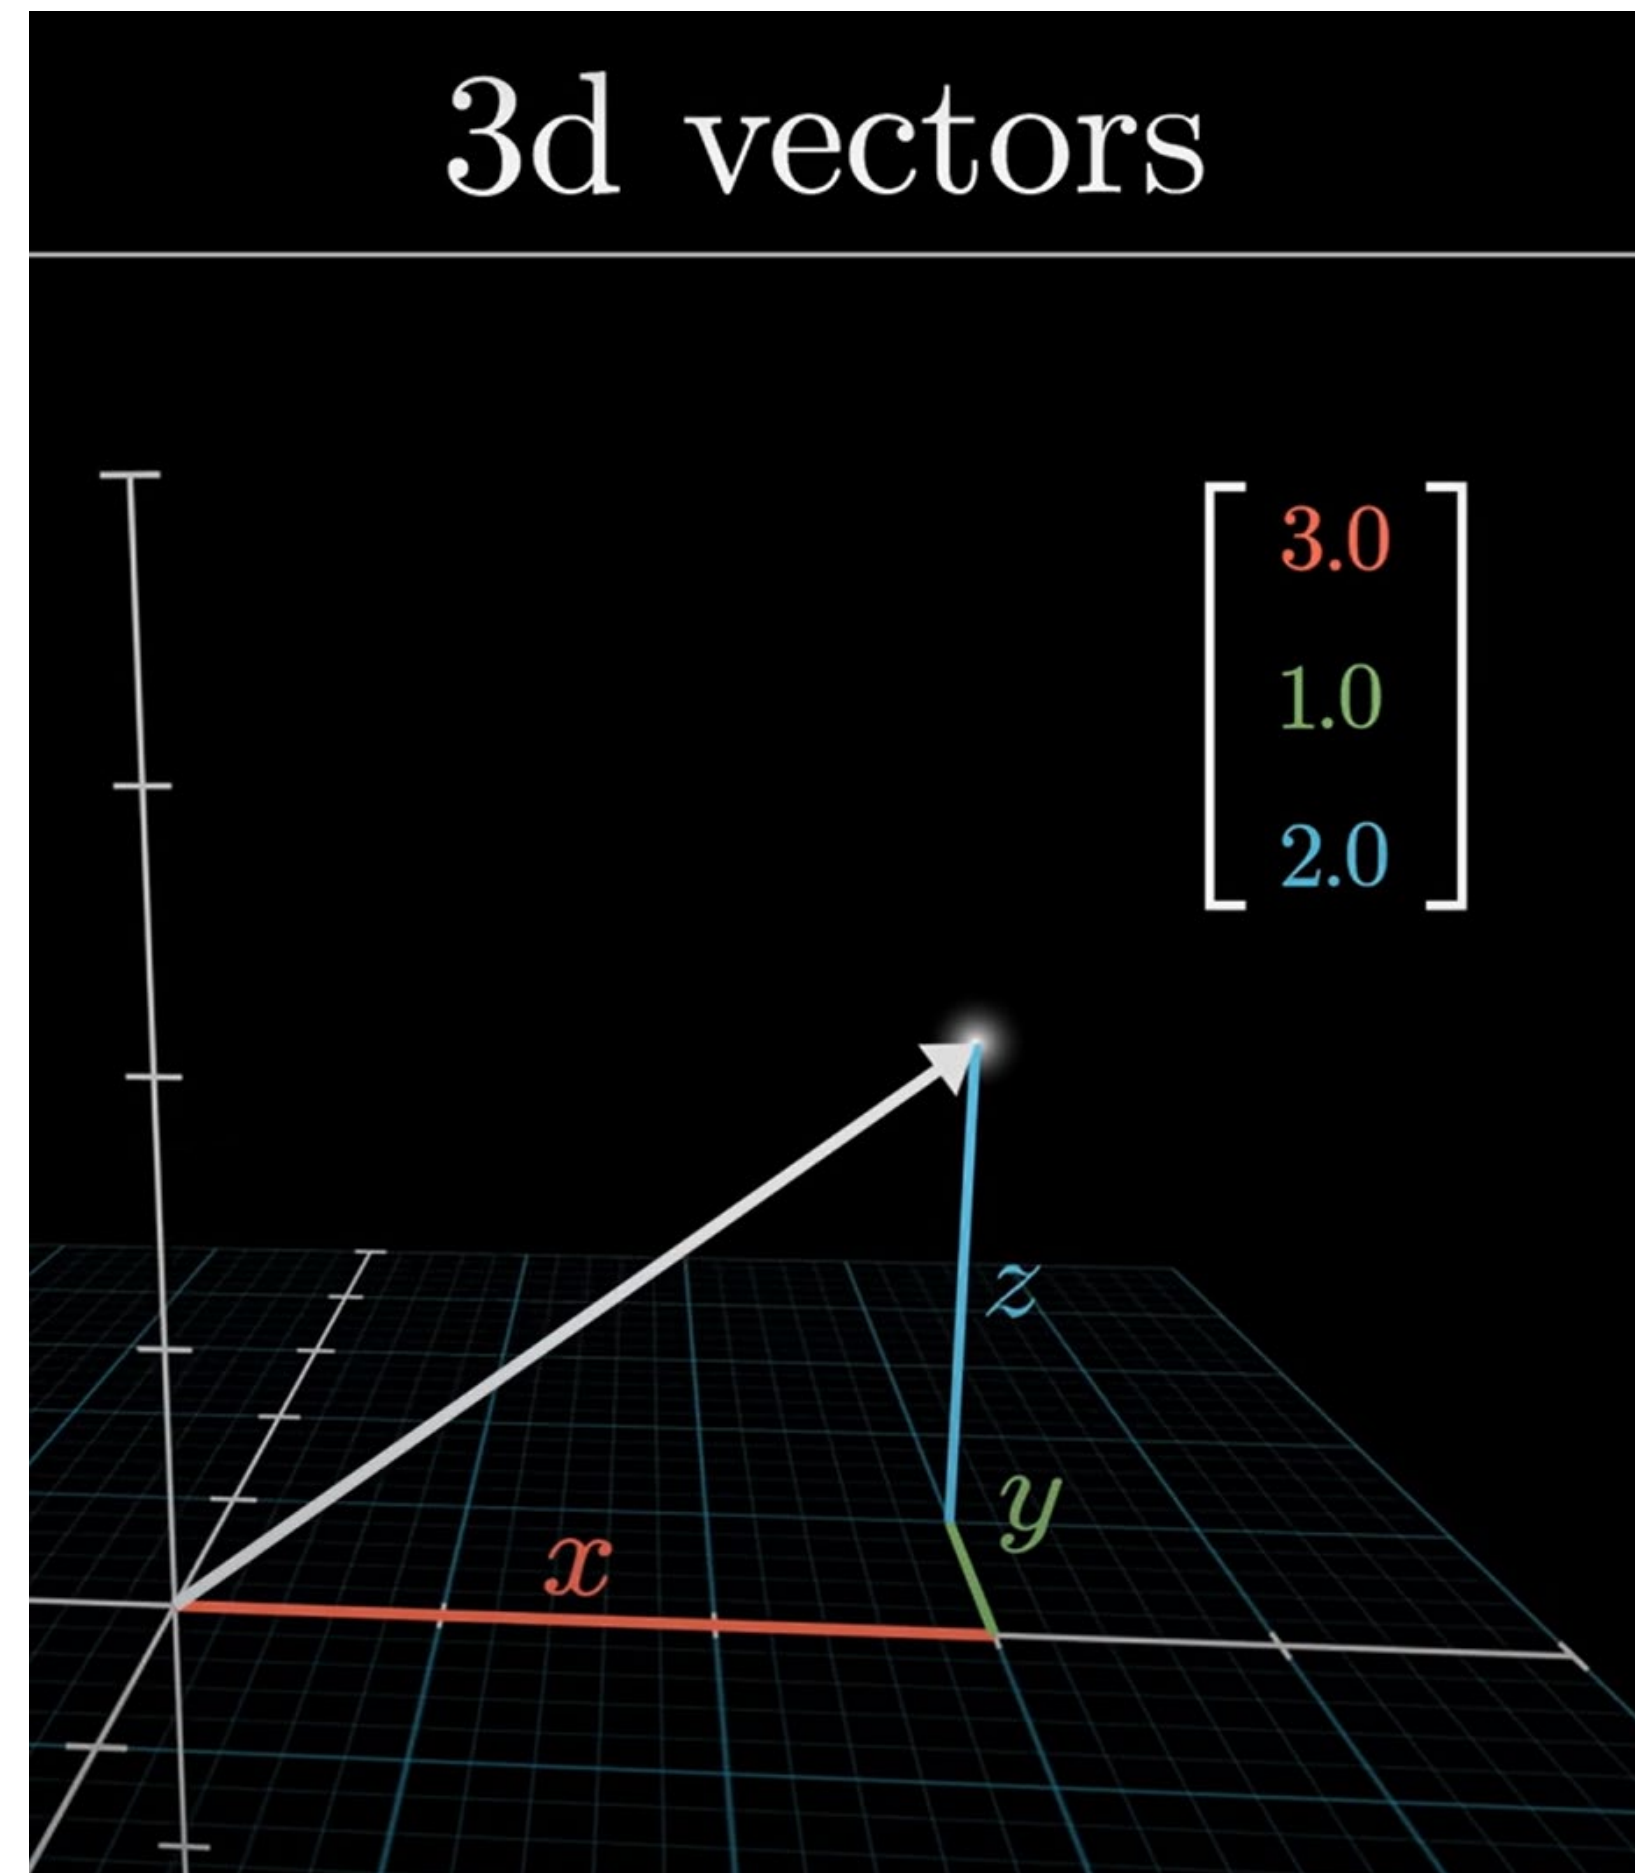
MAKING LLMS USEFUL: RAG

HOW DO YOU FIND THE RIGHT INFORMATION



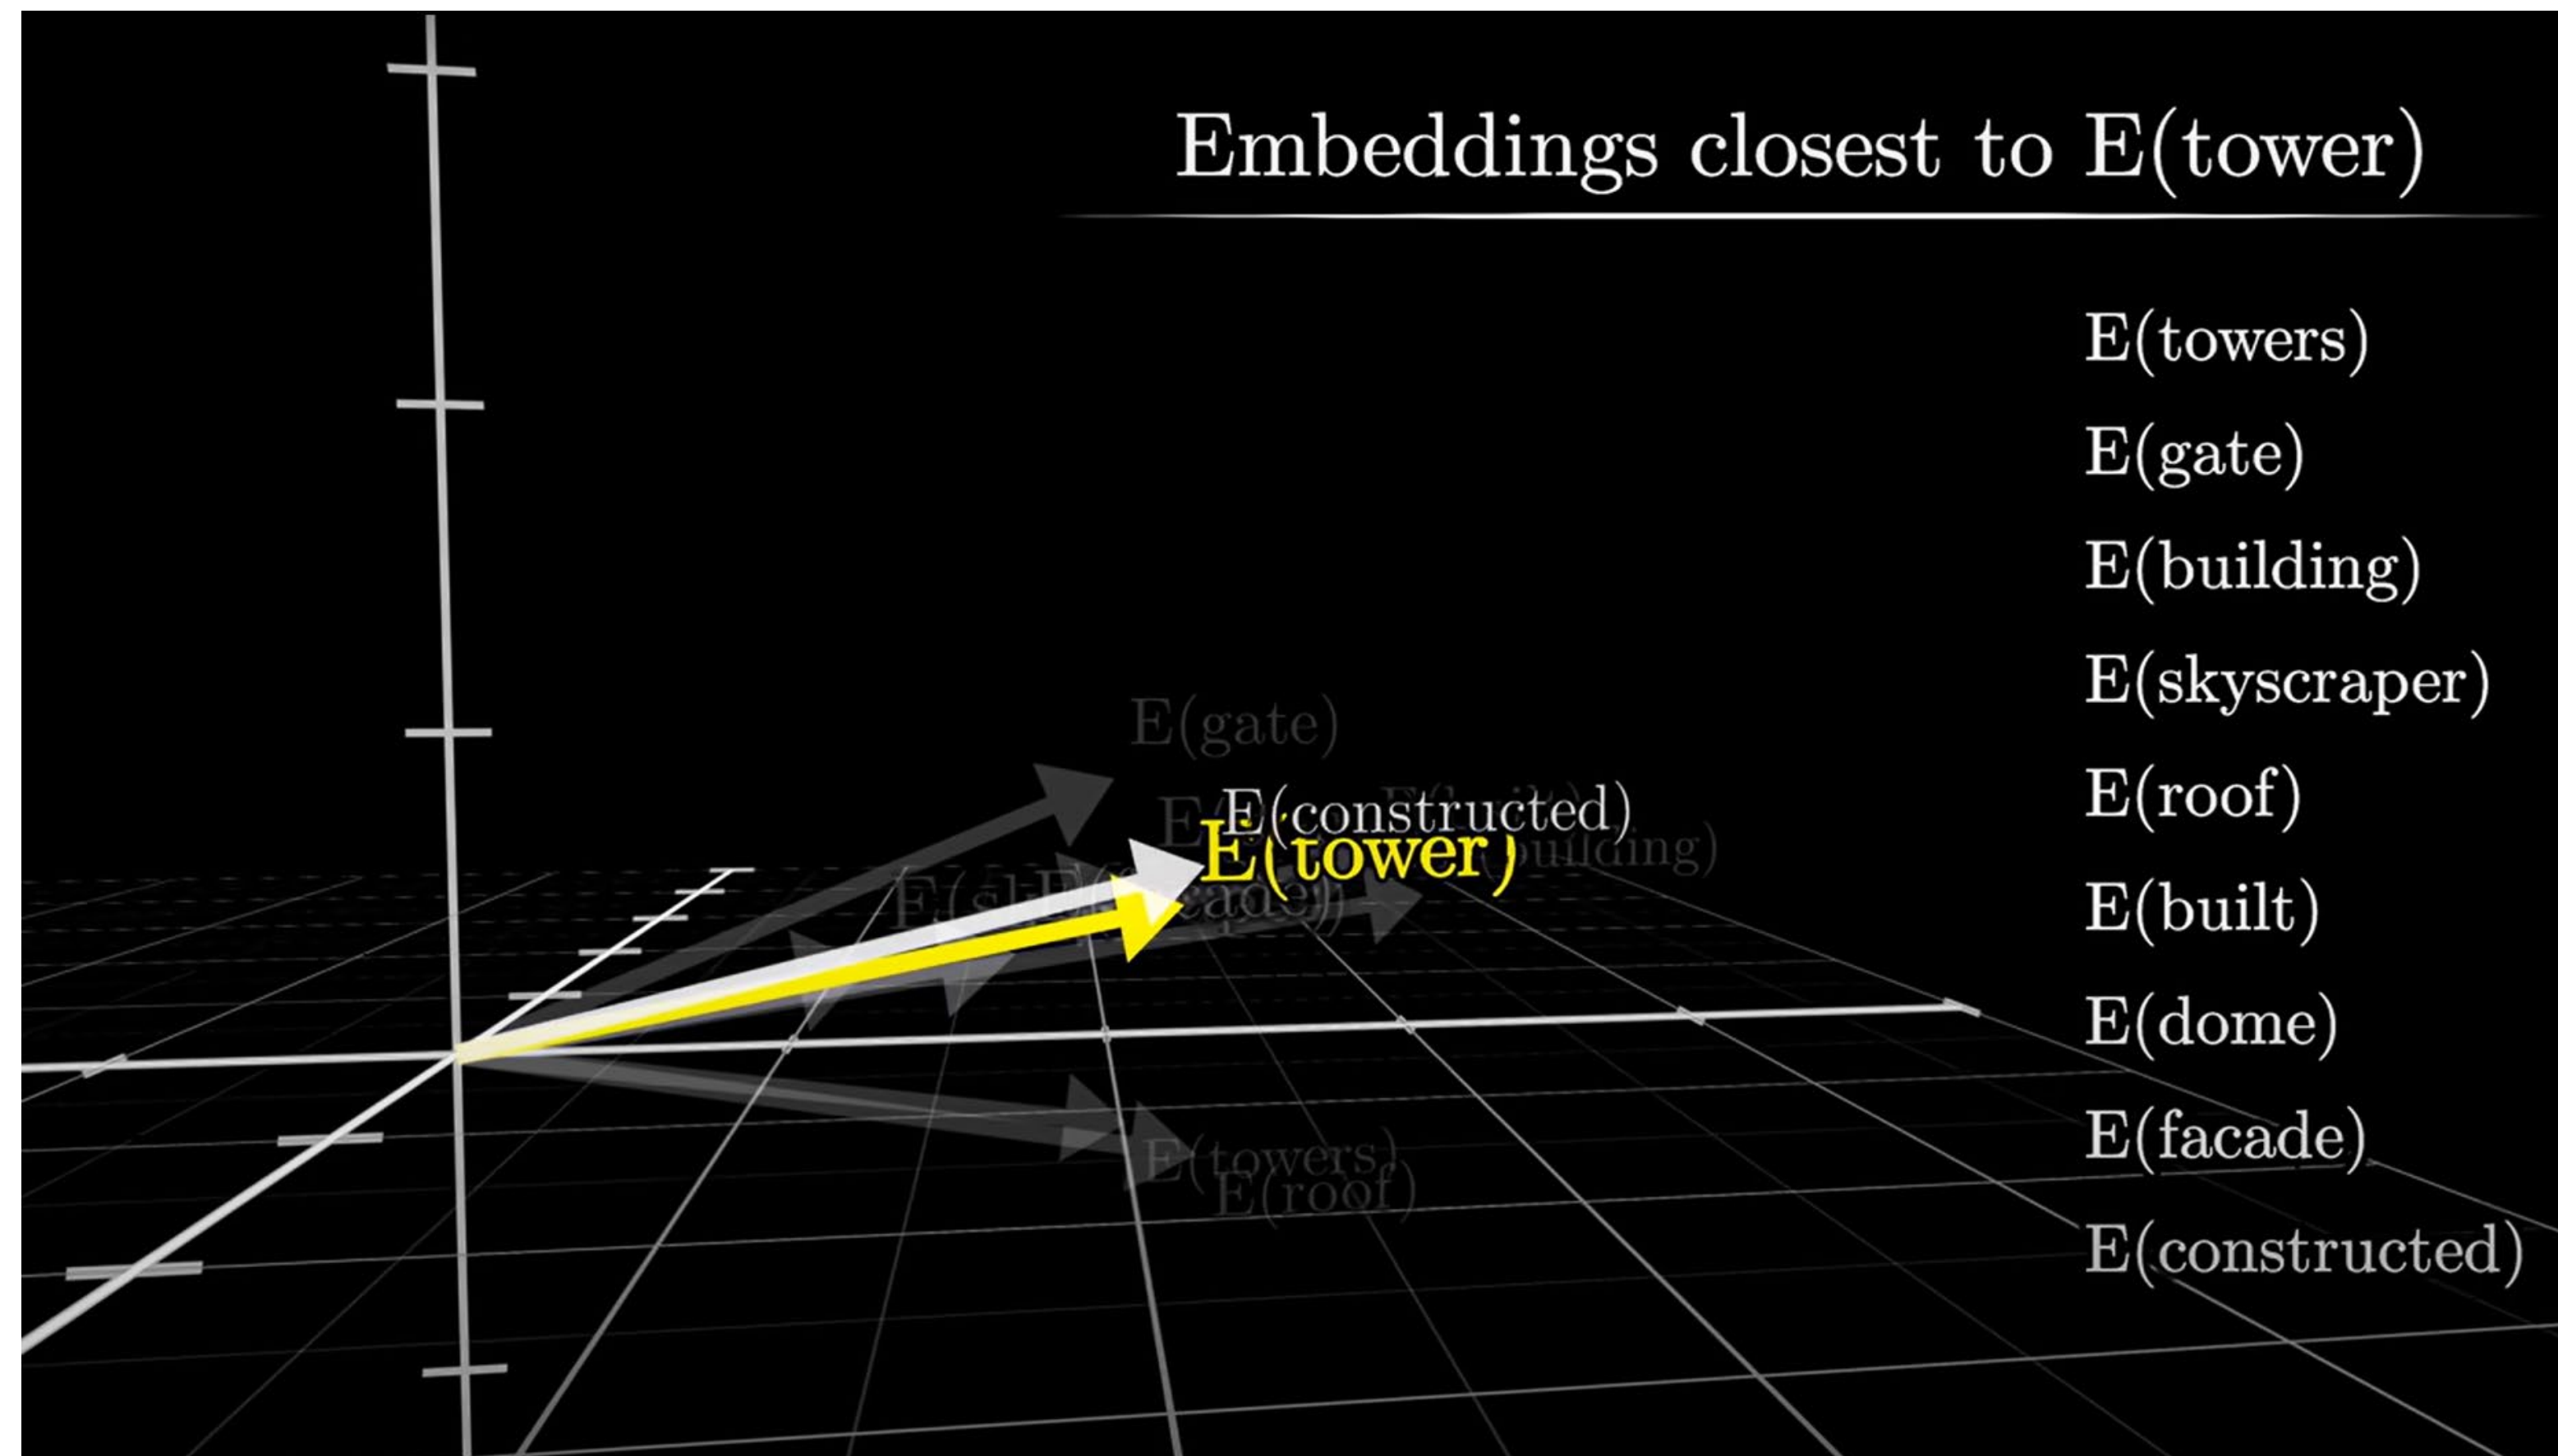
MAKING LLMS USEFUL: RAG

VECTOR EMBEDDINGS



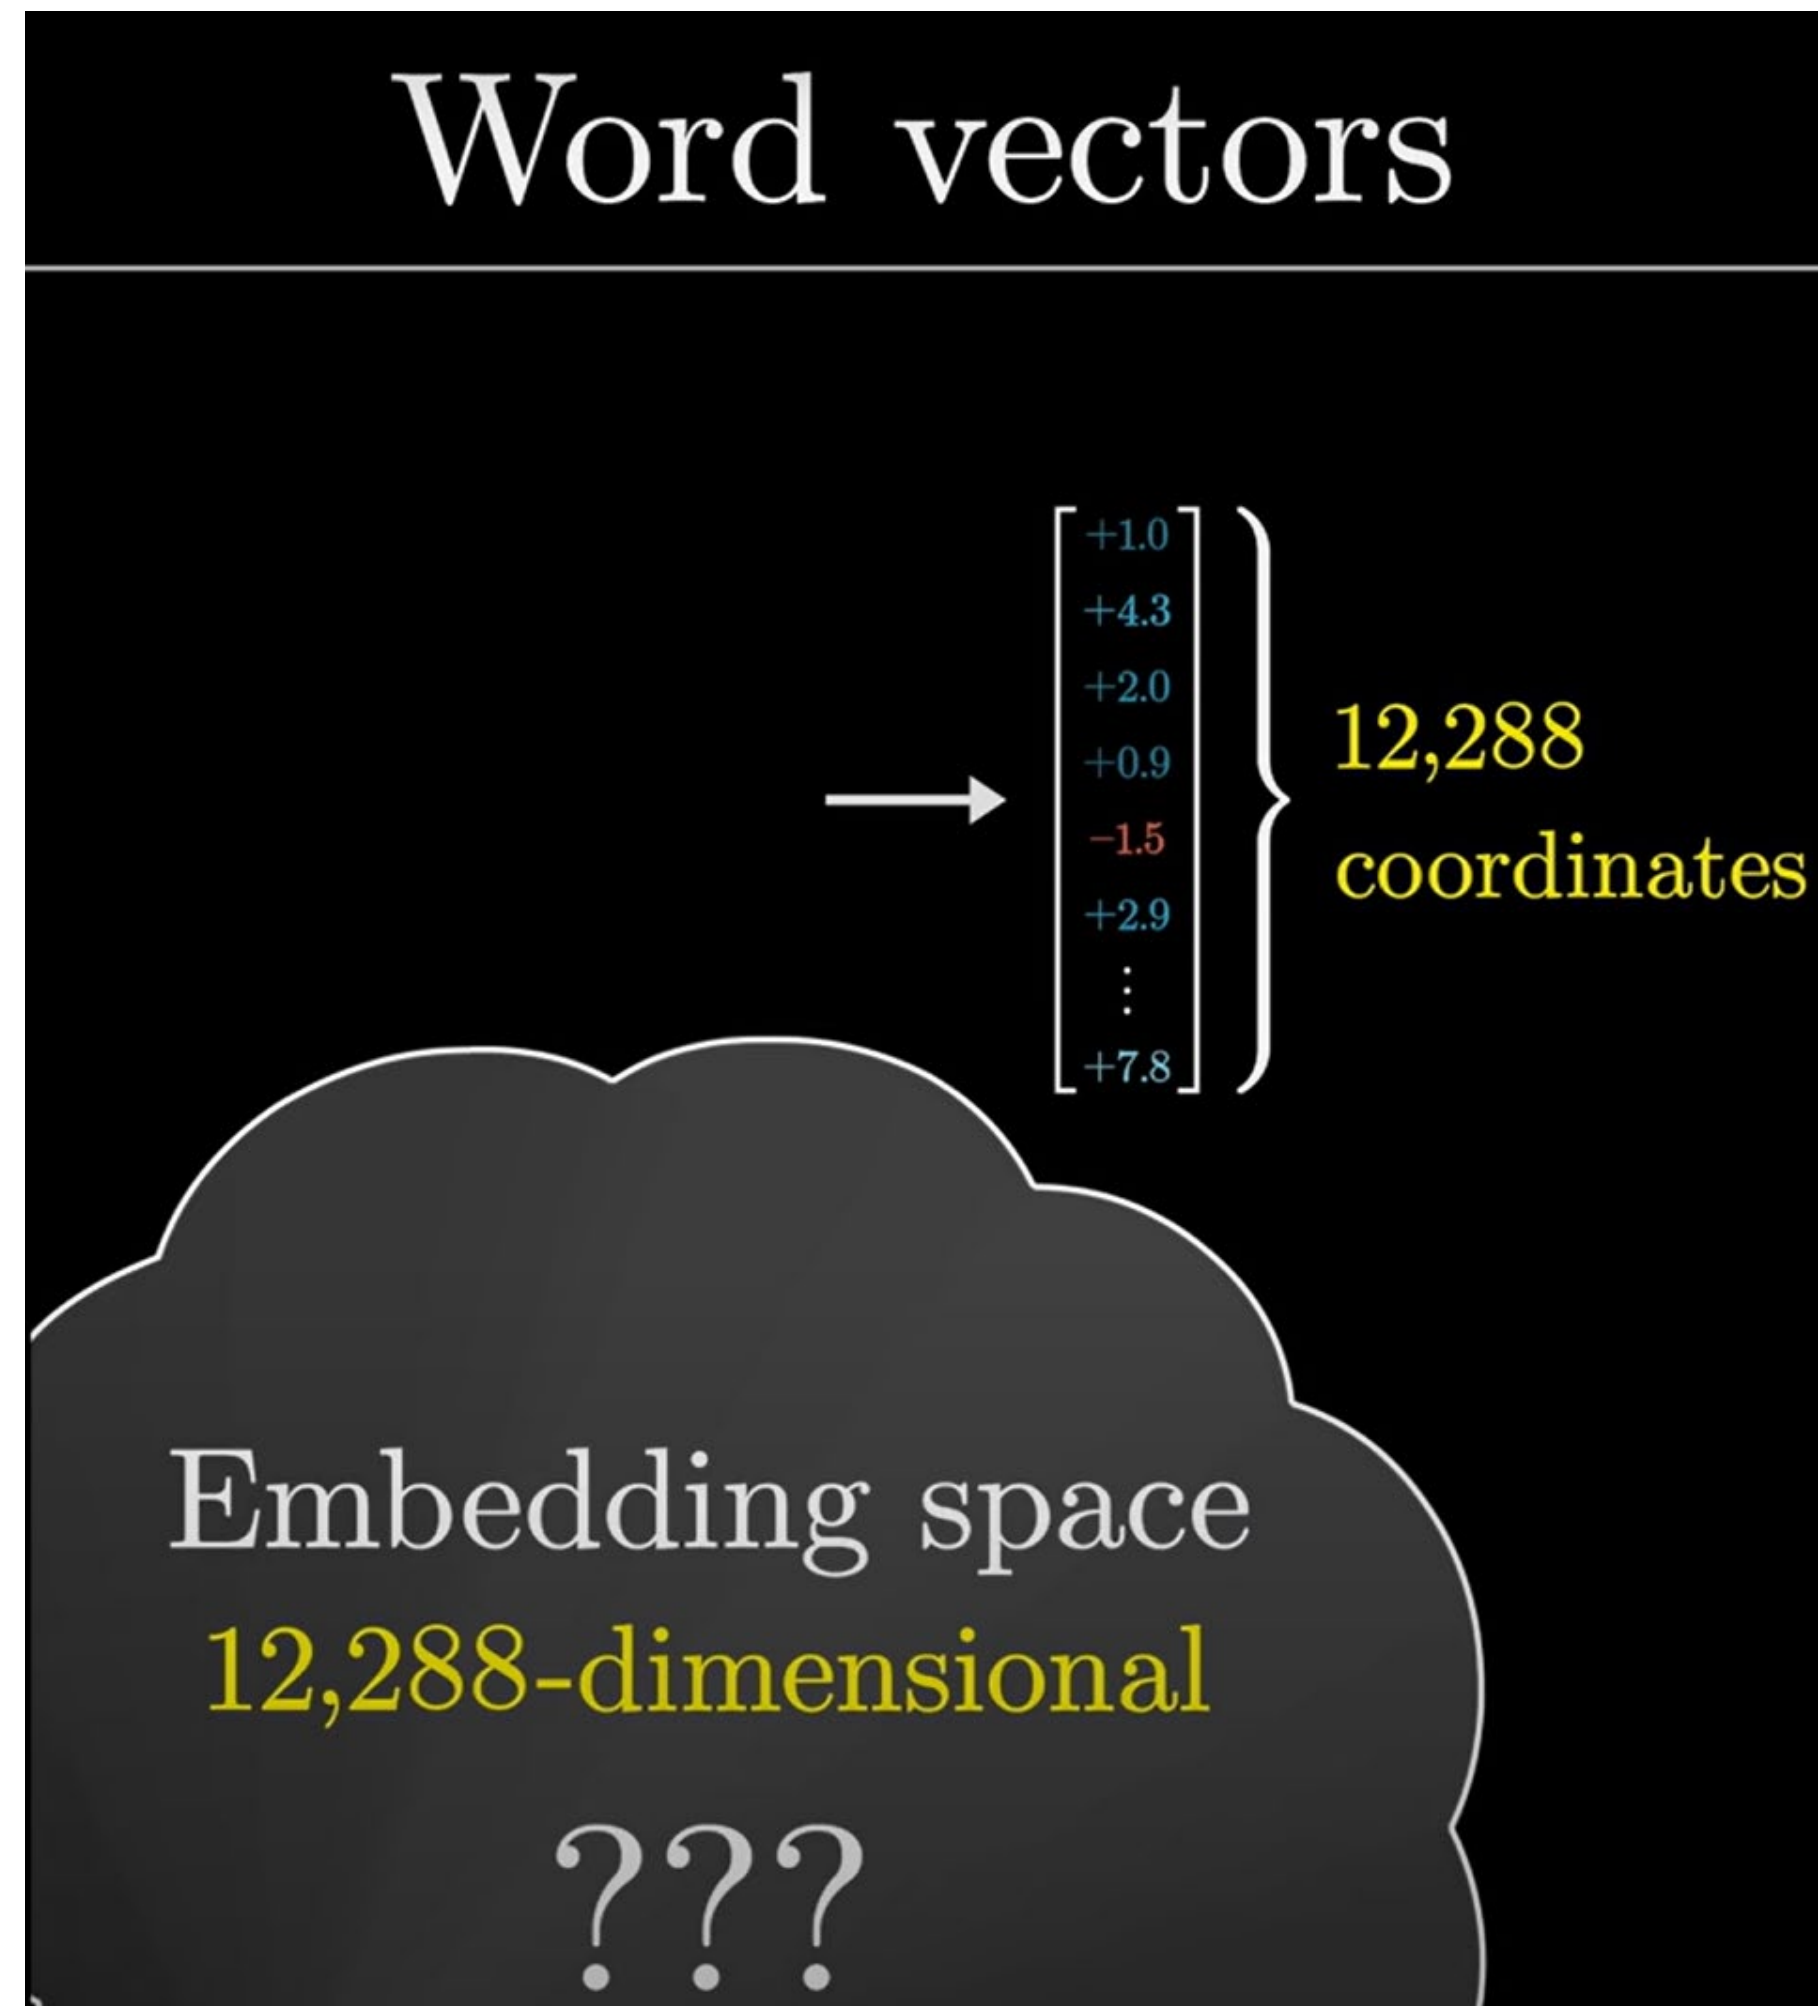
MAKING LLMS USEFUL: RAG

VECTOR EMBEDDINGS



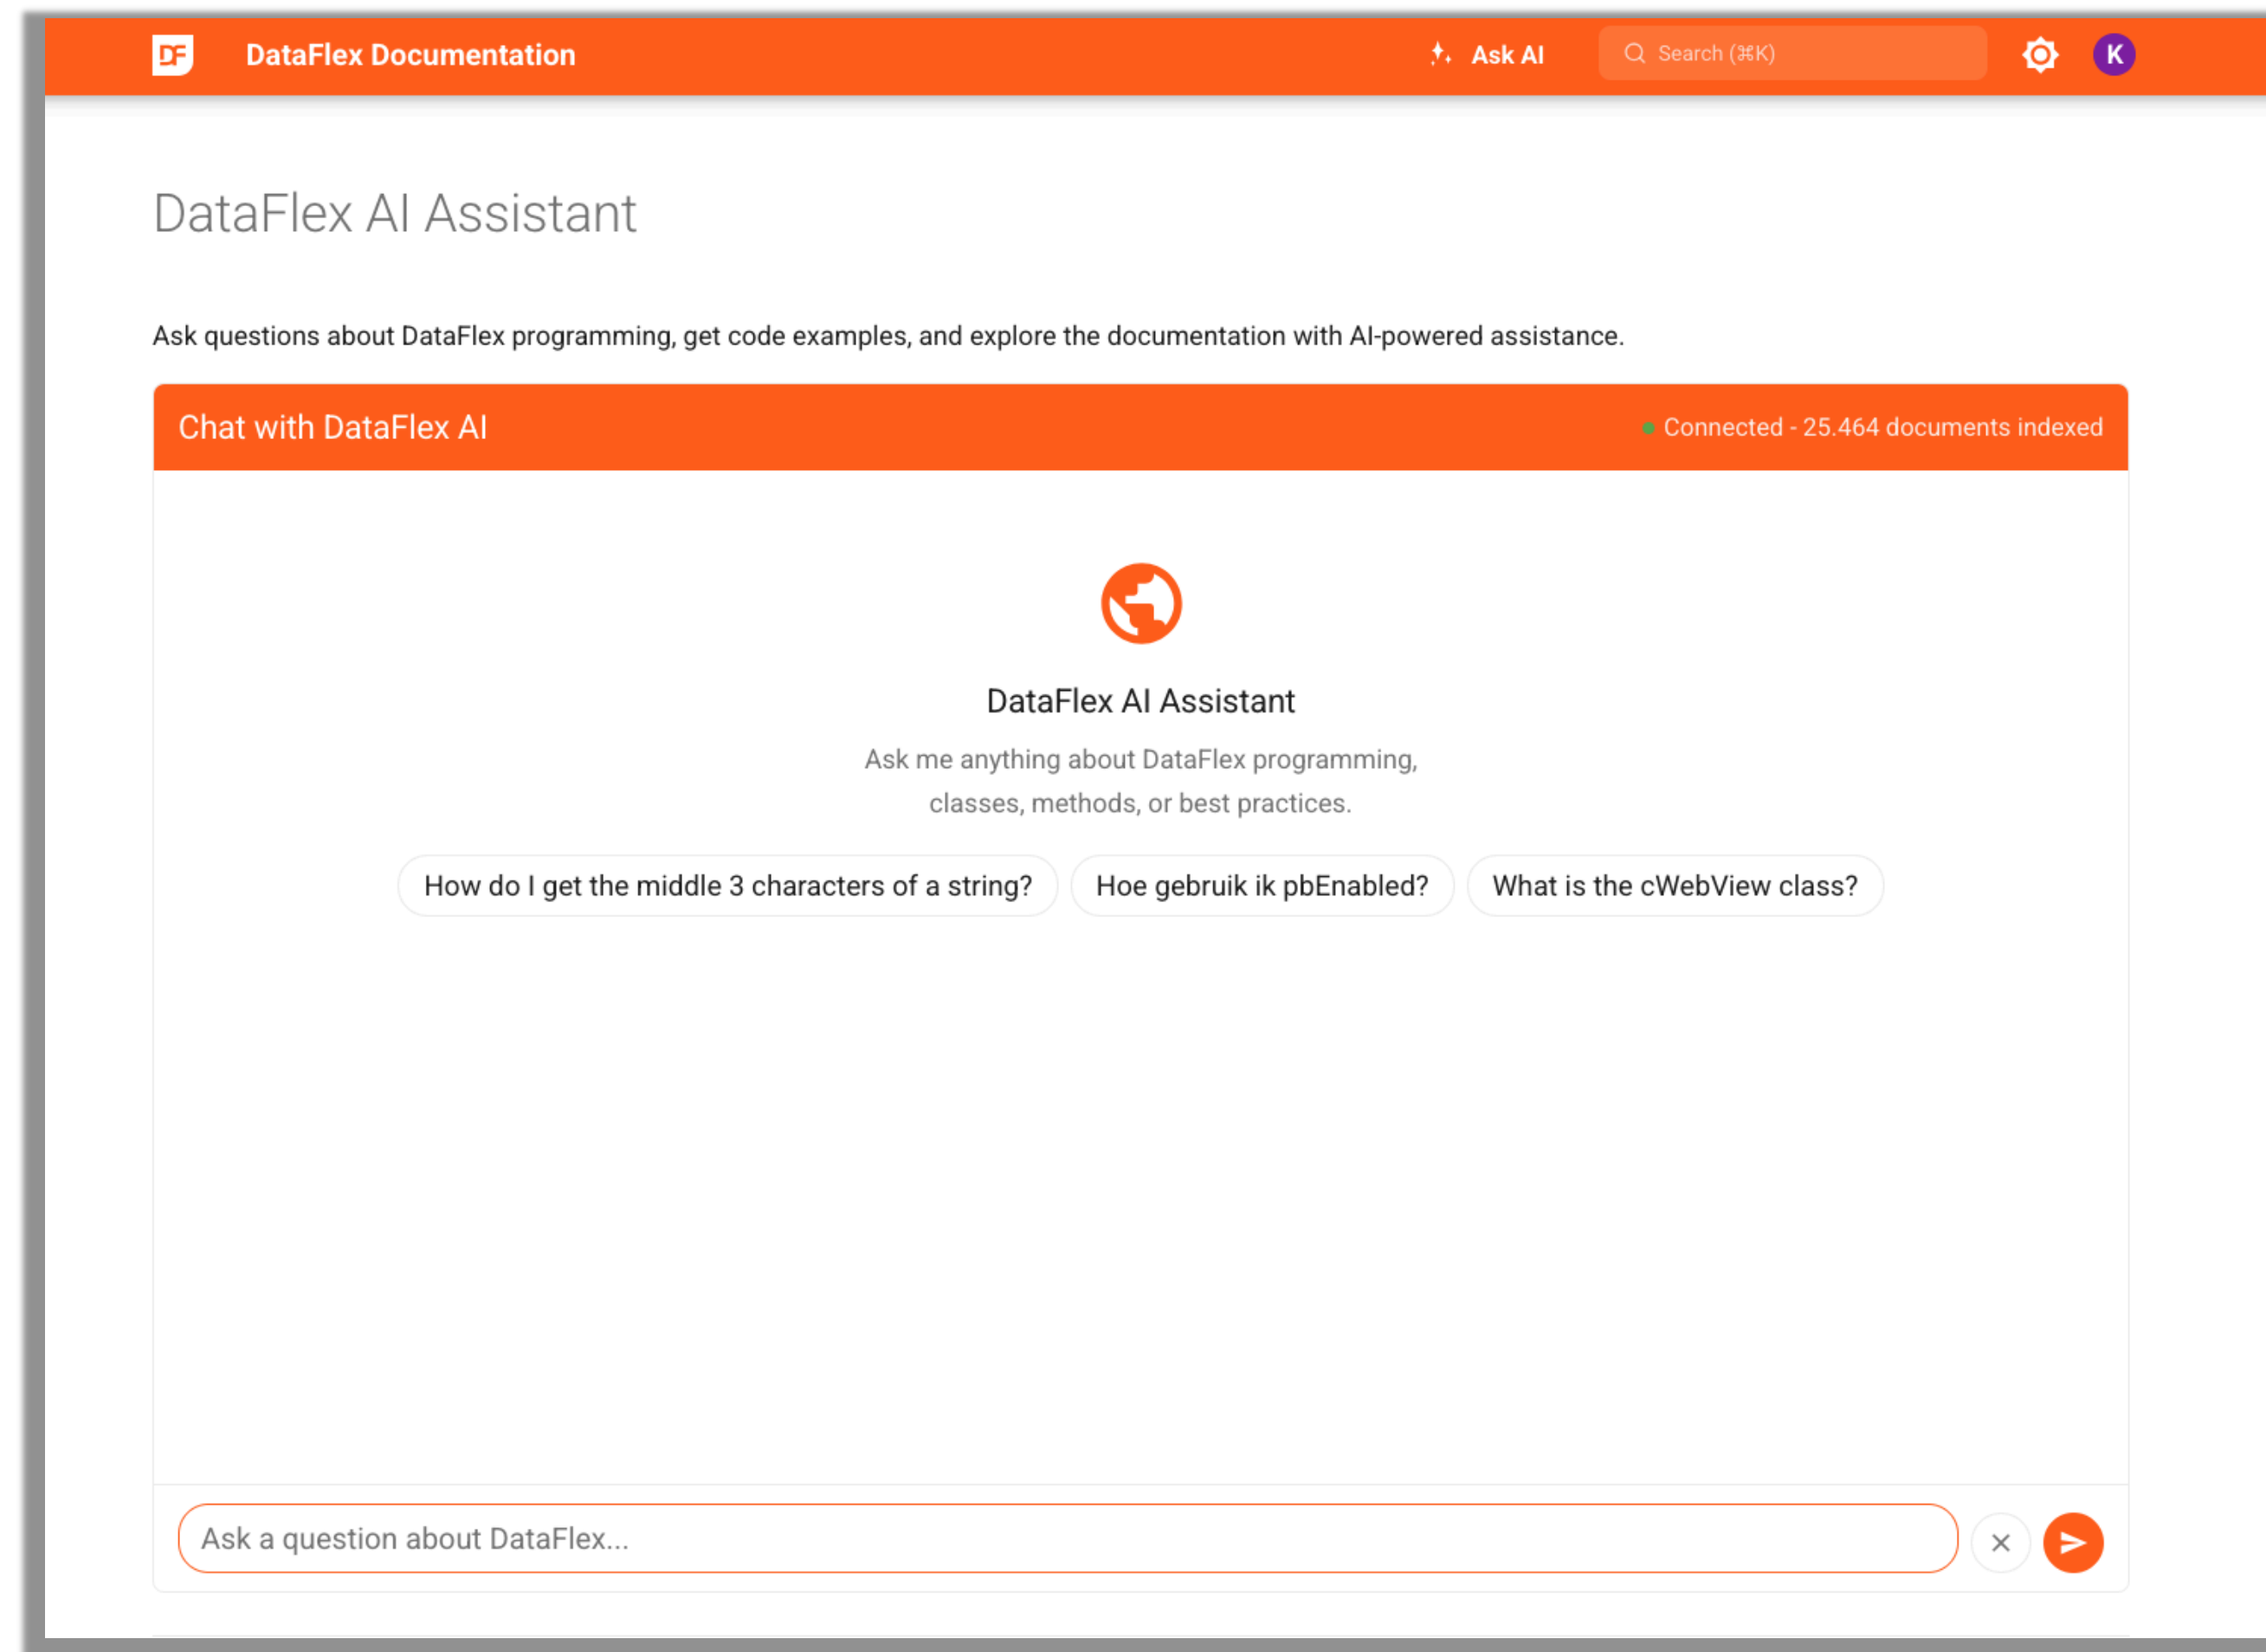
MAKING LLMS USEFUL: RAG

VECTOR EMBEDDINGS

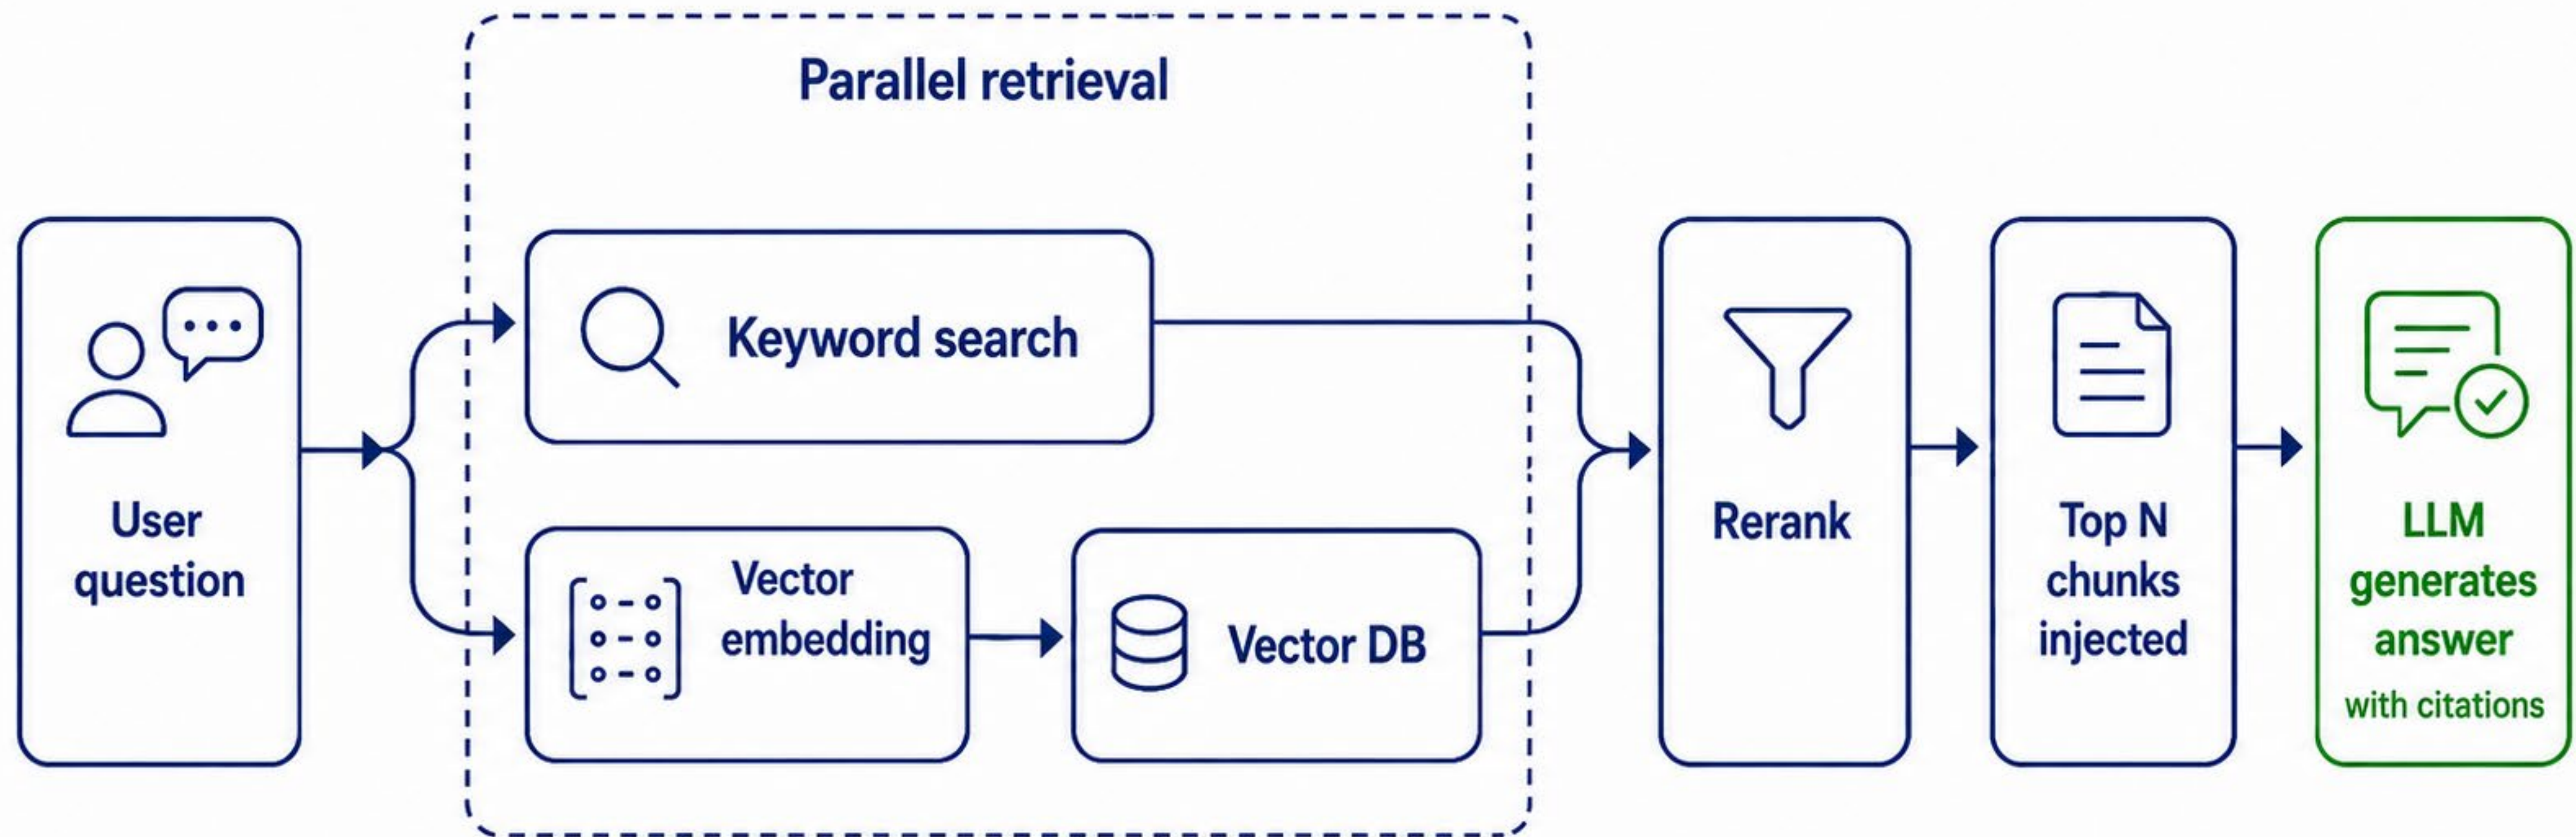


DEMO: RAG-BACKED AI CHAT

AI CHAT LET'S SEE IT

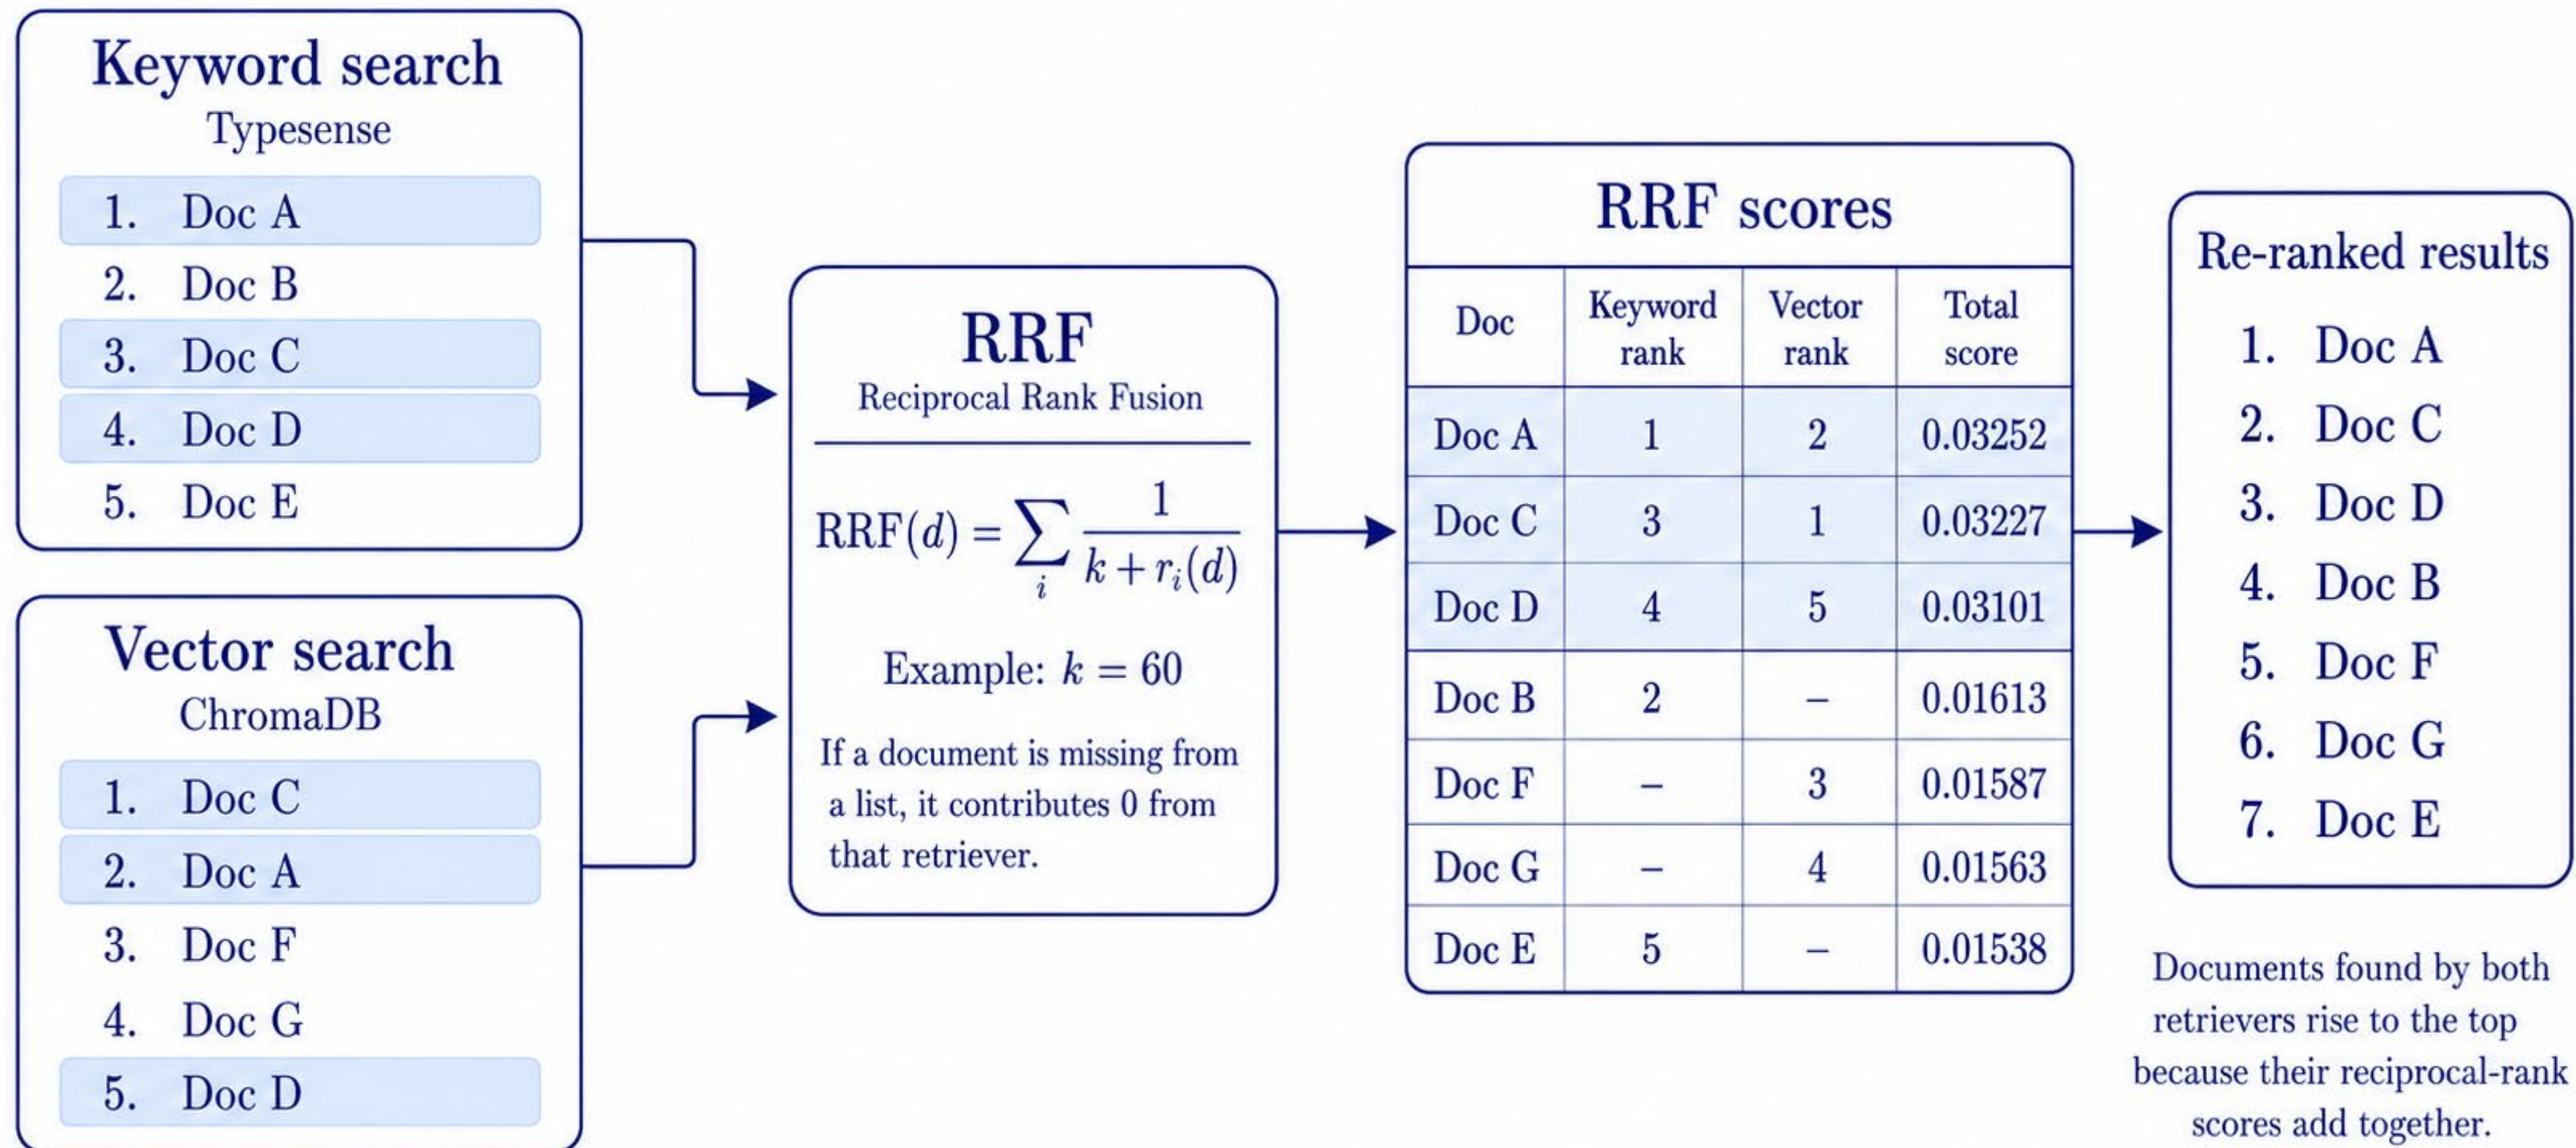


AI CHAT BEHIND THE SCENES



MAKING LLMS USEFUL: RAG

COMBINING KEYWORD AND SEMANTIC SEARCH



AI CHAT

HOW WE KEEP IMPROVING

Prompt And Response

Prompt

How can I build a psMask for a cWebForm so a user can enter a time value like 13:25:01?

Response

To allow a `cWebForm` to display a time like `13:25:01`, set its `psMask` to a time mask and explicitly set the data type so the mask is interpreted correctly. For a 24-hour time with seconds, the mask would be:

```
``dataflex  
Set peDataType to typeAscii  
Set psMask to "hh:mm:ss"  
``
```

The time mask tokens support:

- `h` = hour `0-9`

User Feedback Comment

peDataType should be set to typeTime and not typeAscii or it will not work.

AI CHAT

HOW WE KEEP IMPROVING

Rank	Source File	Section Type	Platform	Rank Score
1	docs/VdfClassRef/Web/cWebBaseDEO-Property-psMask.md	VdfClassRef	Web	3.70553846
2	docs/VdfClassRef/Windows/cCJGridColumn-Property-psMask.md	VdfClassRef	Windows	0.50000000
3	docs/VdfClassRef/Web/cWebColumnDate-Property-psFooterMask.md	VdfClassRef	Web	0.44202898
4	docs/VdfClassRef/Web/cWebColumnSuggestion-Property-psFooterMask.mc	VdfClassRef	Web	0.43571428
5	docs/VdfClassRef/Web/cWebForm-Property-psInputMode.md	VdfClassRef	Web	0.42957746
6	docs/VdfClassRef/Web/cWebColumn-Property-psFooterMask.md	VdfClassRef	Web	0.42361111
7	docs/VdfClassRef/Web/cWebColumnTags-Property-psFooterMask.md	VdfClassRef	Web	0.41216216
8	docs/VdfClassRef/Web/cWebColumnCheckbox-Property-psFooterMask.md	VdfClassRef	Web	0.40666666

AI CHAT

HOW WE KEEP IMPROVING

Description

The psMask property provides support for data masking.

This mask is used for display, and in edit mode if `peDataType` is set to `typeASCII`.

For data types (`peDataType`) other than `typeASCII`, psMask is not used while a control is in edit mode (has the focus). In edit mode, it uses the `psTimeFormat` which can be set in the `cWebApp` object.

You should set `peDataType` explicitly for psMask to work as intended.

Set psMask to manually assign a mask to the text displayed by the control. For more information on valid mask strings, see [Field_Mask](#).

If the control has a data binding (`Entry_Item`), then data masking is automatically set according to the `DataDictionary` rules ([Field_Mask](#)).

Without a data binding, you have to set `piPrecision` and `piMaxLength` for numeric controls.

To use a specific mask (in conjunction with a specific `peDataType` or not), see psMask. If you use `typeNumber` and want to use decimal places, make sure you set `piPrecision`, since it defaults to 0.

AI CHAT

HOW WE KEEP IMPROVING

Description

The psMask property provides support for data masking.

This mask is used for display, and in edit mode if `peDataType` is set to `typeASCII`.

For data types (`peDataType`) other than `typeASCII`, psMask is not used while a control is in edit mode (has the focus). In edit mode, it uses the `psTimeFormat` which can be set in the `cWebApp` object.



AI CHAT

HOW WE KEEP IMPROVING

1. Better docs — clearer text, better examples
2. Better retrieval — clearer chunks, tweak relevance scores
3. Better behaviour — system instructions that tell the AI when to refuse, when to ask back, how to cite.

AI CHAT

HOW WE KEEP IMPROVING

- Less mistakes through feedback
- Increase knowledge:
 - DataFlex books
 - DataFlex Learning Center videos
 - DataFlex Reports

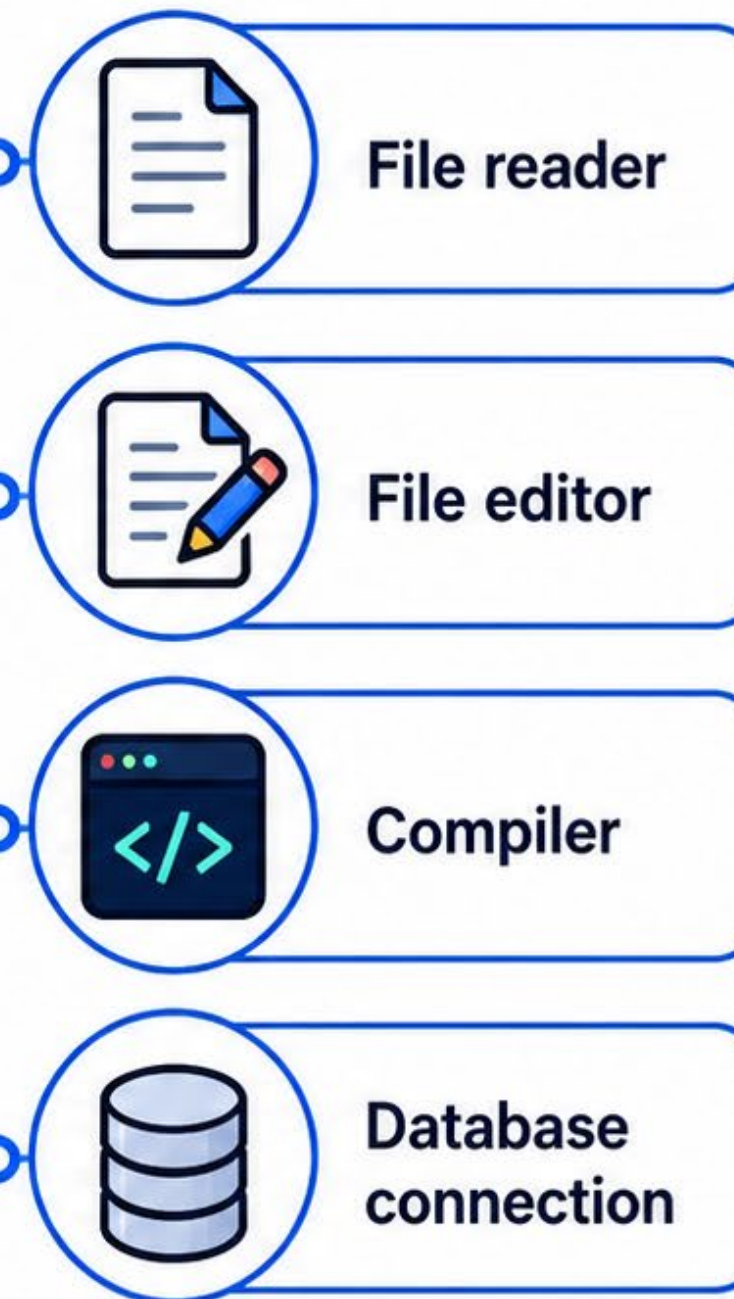
MAKING LLM'S USEFUL: MCP

MAKING LLMS USEFUL: MCP MODEL CONTEXT PROTOCOL

Different AI assistants



Available tools



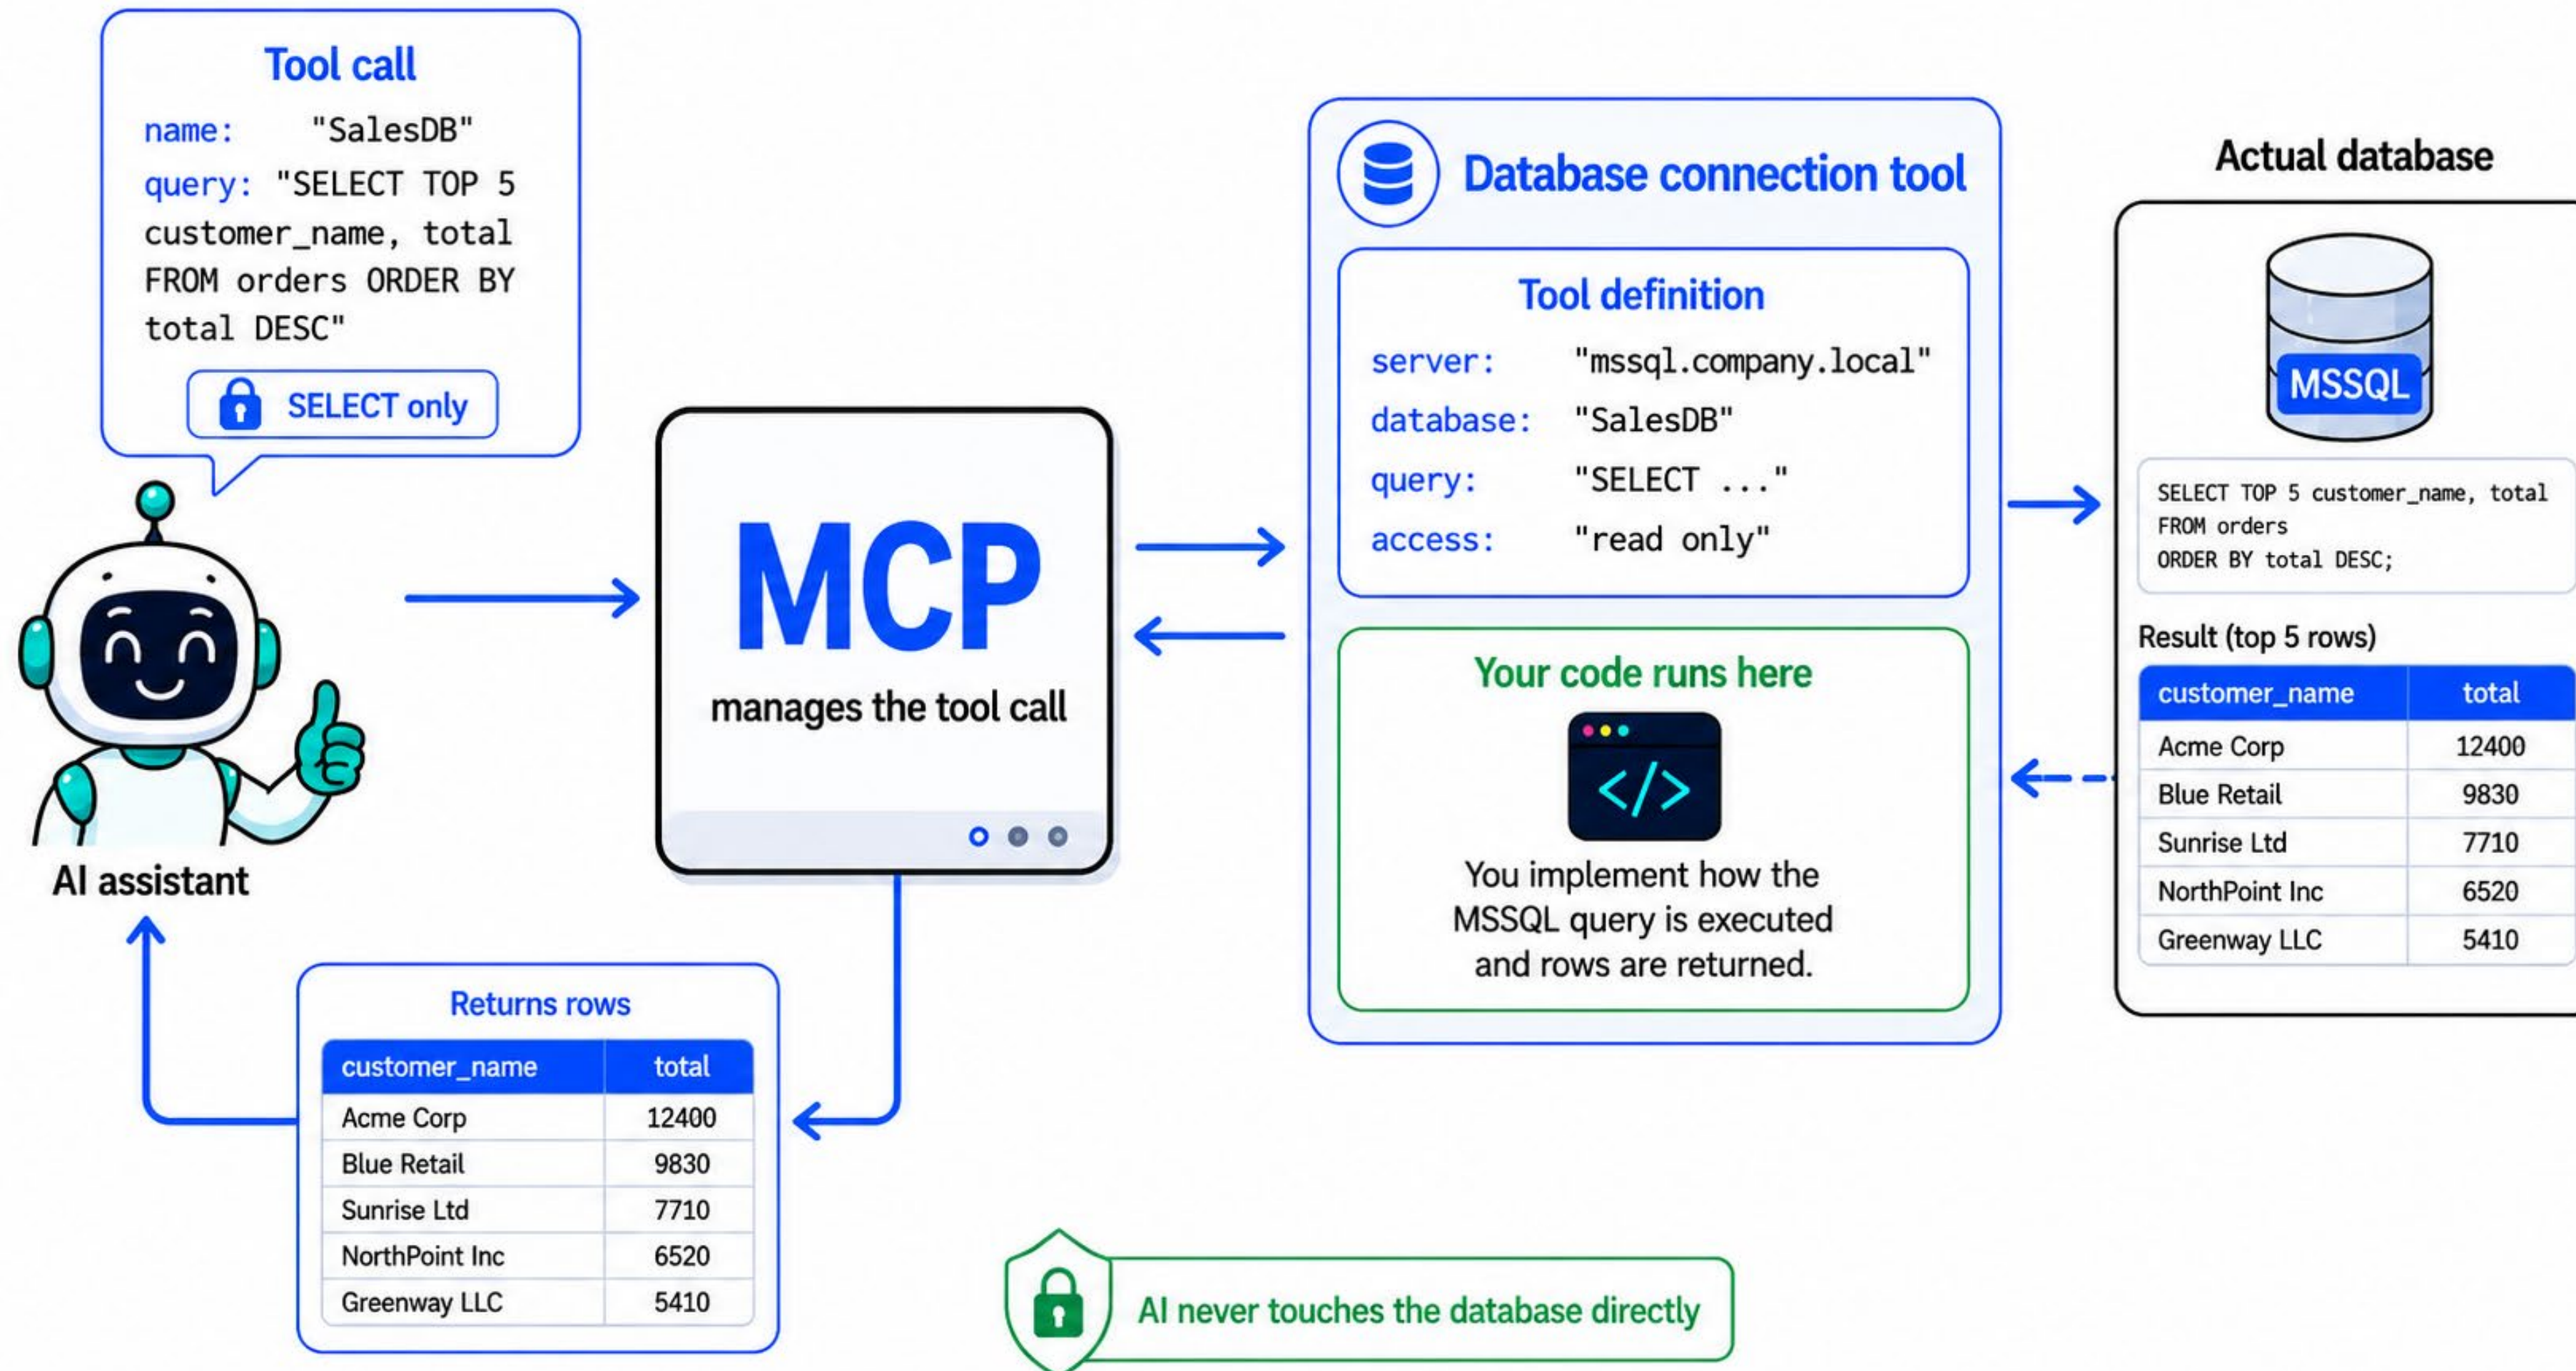
Many AI assistants ↔ one standard MCP ↔ many tools

MAKING LLMS USEFUL: MCP

THE CONTROL LAYER

- Responsibilities of the MCP layer:
 - Lists the tools
 - Holds each tool's definition
 - Describes behavior
 - Enforces the rules
 - Routes and returns

MAKING LLMS USEFUL: MCP TOOLS



DEMO: AI AWARE SQL DATABASES

AI AWARE SQL DATABASES

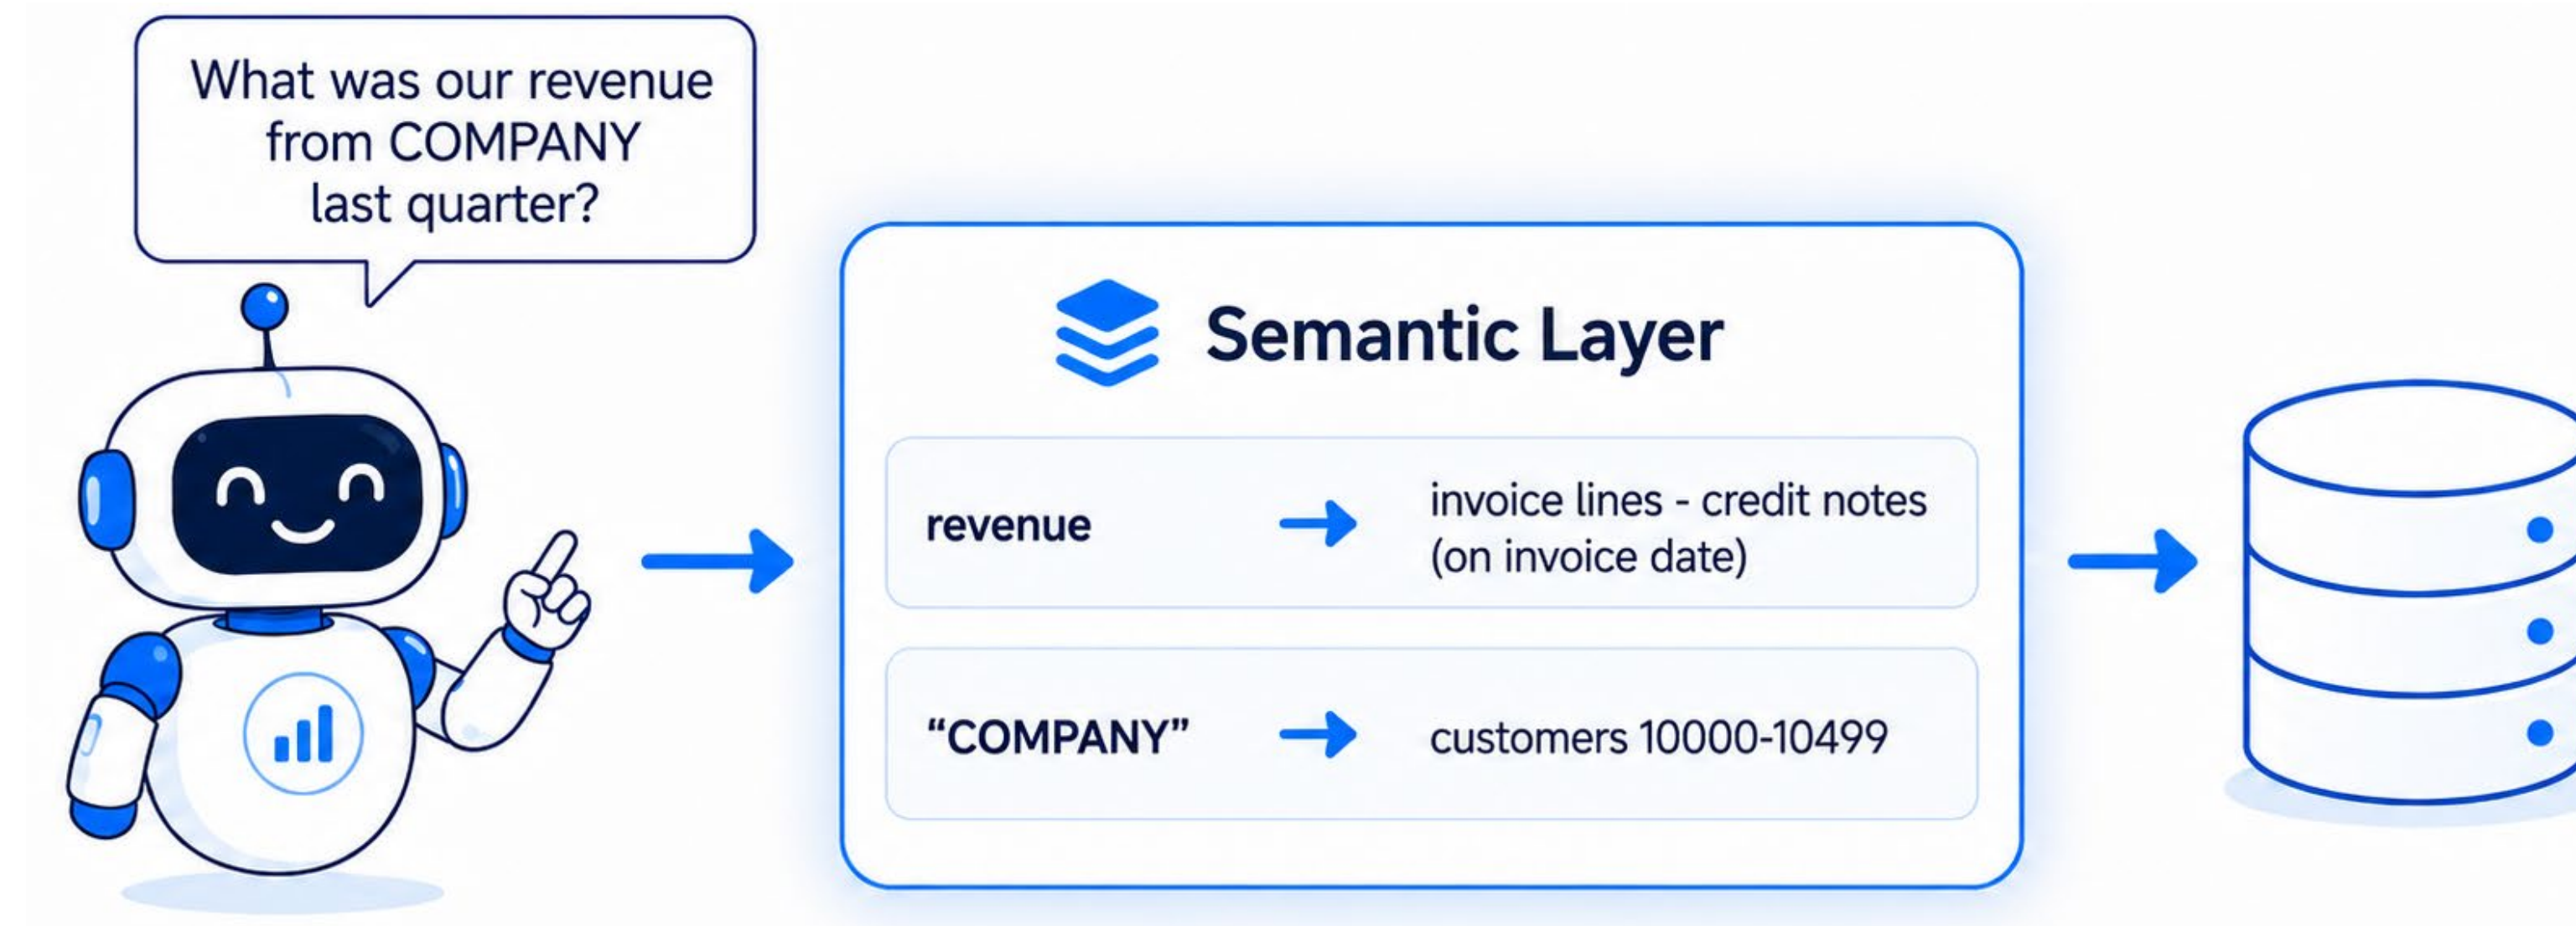
AI IN SQL DATABASES

- Early research and development phase
- End-user focused
- You decide data input, now also output

AI AWARE SQL DATABASES

THE SEMANTIC LAYER

- Generic tool calling: AI writes raw SQL, guessing
- With semantic layer: AI works with **business concepts you defined**
- You control the meaning → correct, consistent, governed answers



AI AWARE SQL DATABASES

EXAMPLES – ADVENTUREWORKS2012

- “What is our total revenue?”

What the AI picks	Result
SUM(SalesOrderHeader.TotalDue) — includes tax + freight	\$123,216,786
SUM(SalesOrderHeader.SubTotal)	\$109,846,381
SUM(SalesOrderDetail.LineTotal) ← defined revenue	\$109,846,381

- \$13.4M (12%) overstated just by grabbing obvious column

AI AWARE SQL DATABASES

EXAMPLES – ADVENTUREWORKS2012

- “How many orders did we have?”

What the AI picks	Result
COUNT(*) on the joined rows	121,317
COUNT(DISTINCT SalesOrderID) ← defined order_amount	31,465

- Nearly 4x too high

AI AWARE SQL DATABASES

EXAMPLES – ADVENTUREWORKS2012

- “What is our average order value?”
- AOV = revenue ÷ orders, so errors #1 and #2 stack:
 - Generic AI (TotalDue ÷ line-item count): $\$123.2\text{M} / 121,317 = \sim\$1,016$
 - Defined avg_order (LineTotal ÷ distinct orders): $\$109.8\text{M} / 31,465 = \sim\$3,491$

AI AWARE SQL DATABASES DEMO

DataFlex Insights
Tenant: demo

Road-250 Black, 44	€ 1.280.263,04
Road-250 Black, 48	€ 1.169.567,70
Road-250 Black, 52	€ 1.004.373,92
Touring-1000 Blue, 60	€ 885.952,36

Het best verkopende item van **2007** was dus **Mountain-200 Black, 38** met **€ 2.176.123,80** omzet.

[Toelichting bij dit antwoord](#)

Jij
Can you make a graph of the total revenue per year that you have data of

Insights
Hier is een eenvoudige grafiek van de **totale omzet per jaar** voor alle jaren waar data beschikbaar is. Omzet is geboekte verkoop op orderdatum, **inclusief korting**.

Omzet per jaar

2005	€ 11.331.808,96
2006	€ 30.674.773,18
2007	€ 42.011.037,18
2008	€ 25.828.762,10

Jaar	Omzet
2005	€ 11.331.808,96
2006	€ 30.674.773,18
2007	€ 42.011.037,18
2008	€ 25.828.762,10

2007 had de hoogste omzet met **€ 42.011.037,18**.

[Toelichting bij dit antwoord](#)

Stel je vraag...

Verstuur

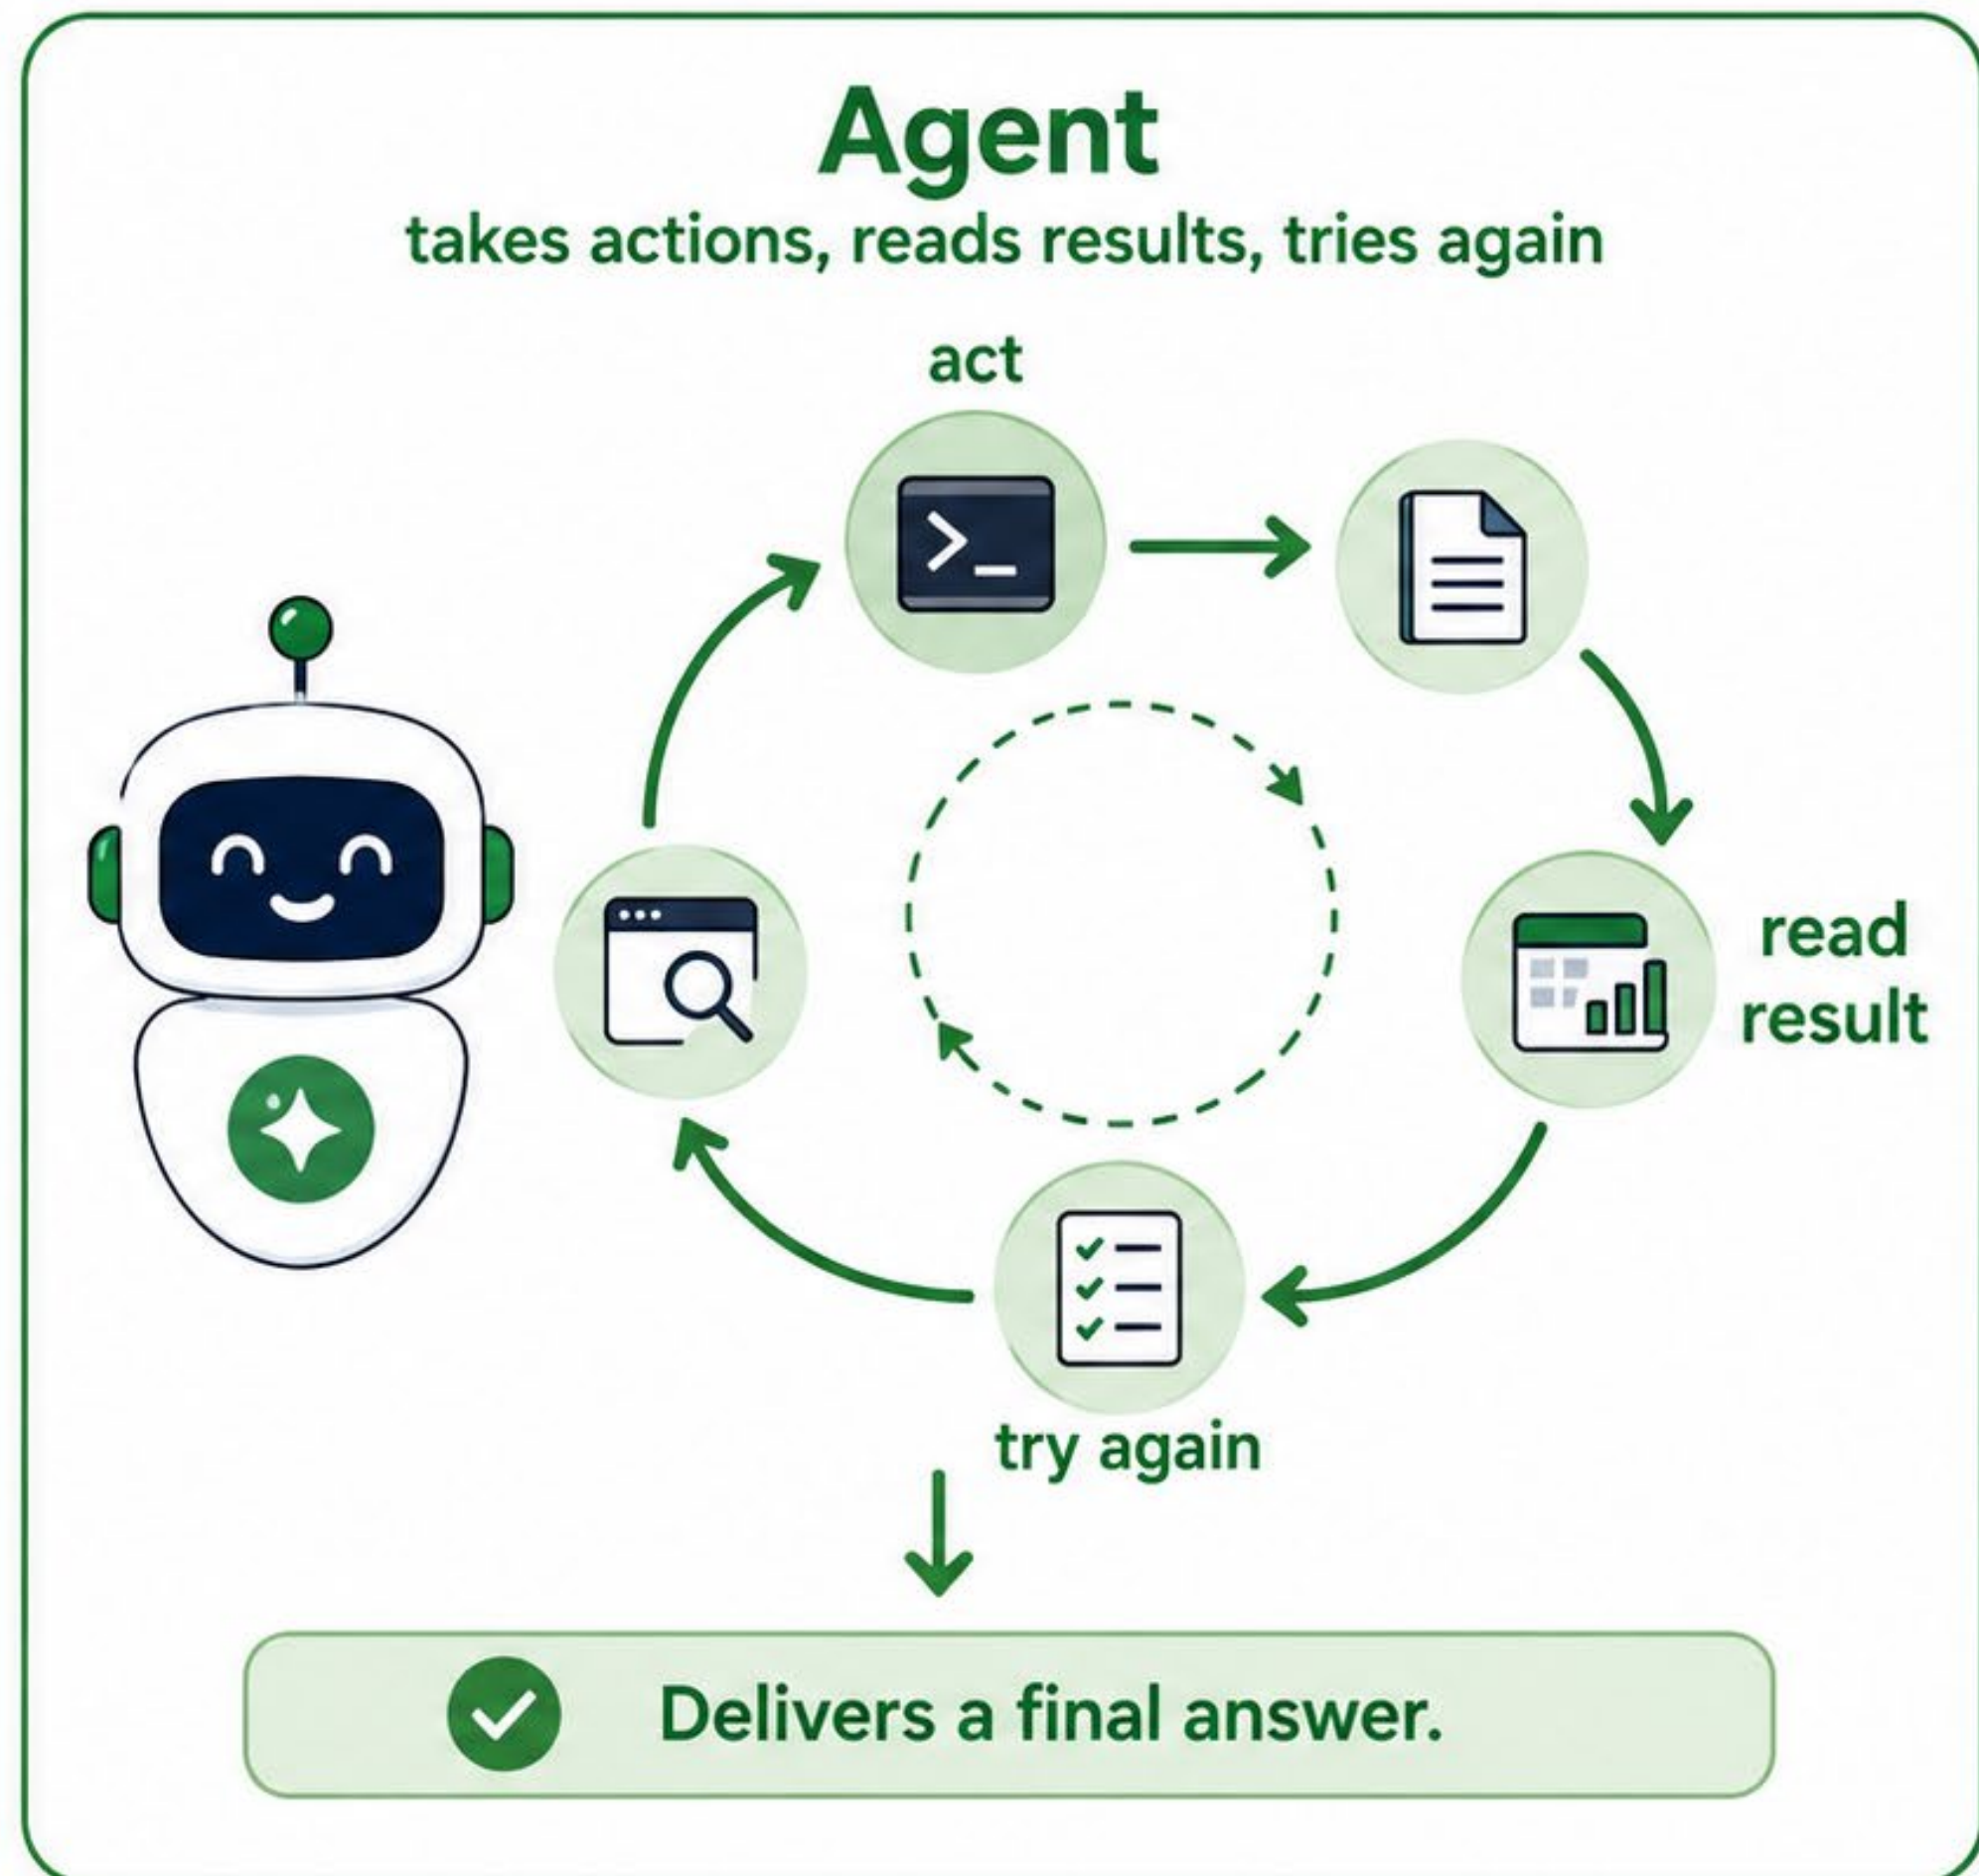
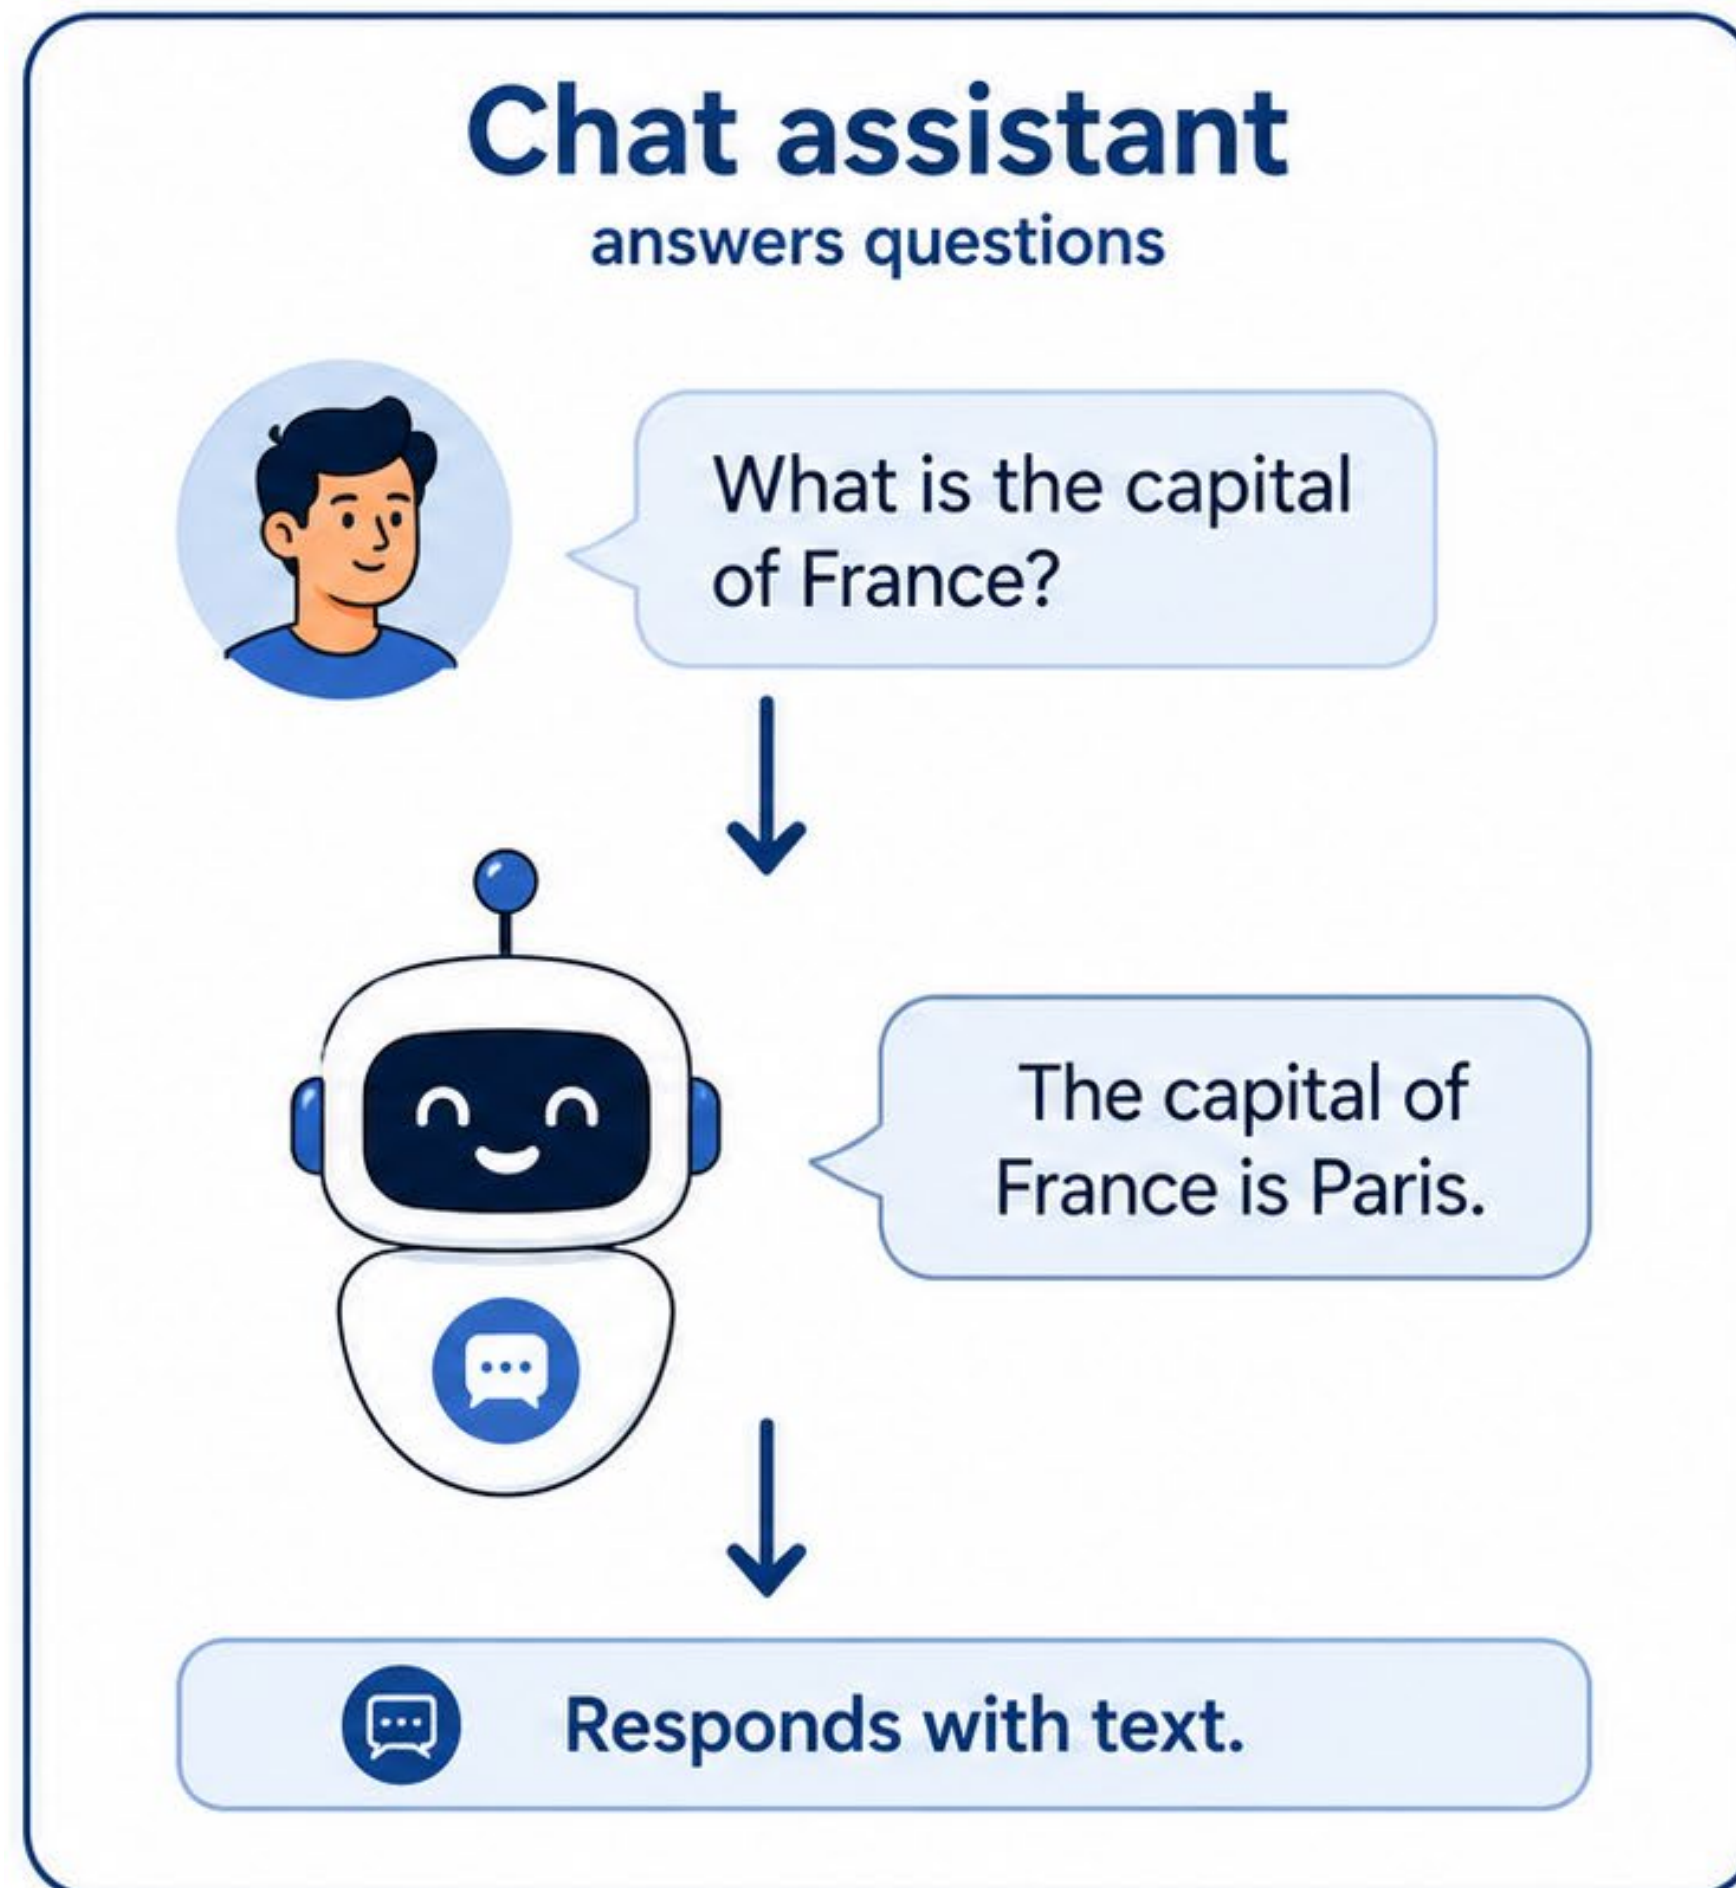
AI AWARE SQL DATABASES WHERE THIS IS HEADED

- Extract from Data Dictionaires -> structural layer
- AI Portal -> judgement layer
- Charts, dashboards, reports

/*WITHOUT A SEMANTIC LAYER, THE AI SPEAKS SQL. WITH IT, THE AI SPEAKS YOUR BUSINESS LANGUAGE*/

DEMO: THE AI AGENT INSIDE THE STUDIO

DATAFLEX STUDIO AI AGENT FROM “AI THAT ANSWERS” TO “AI THAT DOES”



DATAFLEX STUDIO AI AGENT

THE STUDIO AGENT: WHAT IT IS

The screenshot displays the DataFlex Studio 99.0 interface. The main editor shows code for defining web columns. A new column, `oColTotalRevenue`, is highlighted in green. The code defines its entry item as `Customer.Purchases`, sets its caption to "Total Revenue", width to 130, and alignment to `alignRight`. The AI Agent window on the right shows a conversation where the agent provides instructions and a summary of the new column's configuration.

```
55  
56 Object oColCity is a cWebColumn  
57     Entry_Item Customer.City  
58     Set psCaption to "City"  
59     Set piWidth to 160  
60 End_Object  
61  
62 Object oColState is a cWebColumn  
63     Entry_Item Customer.State  
64     Set psCaption to "State"  
65     Set piWidth to 80  
66 End_Object  
67  
68 Object oColPhone is a cWebColumn  
69     Entry_Item Customer.Phone_Number  
70     Set psCaption to "Phone"  
71     Set piWidth to 140  
72 End_Object  
73  
74 Object oColTotalRevenue is a cWebColumn  
75     Entry_Item Customer.Purchases  
76     Set psCaption to "Total Revenue"  
77     Set piWidth to 130  
78     Set pbPromptButton to False  
79     Set peAlign to alignRight  
80 End_Object  
81  
82 Object oColStatus is a cWebColumnCheckbox  
83     Entry_Item Customer.Status  
84     Set psCaption to "Active"  
85     Set piWidth to 70  
86 End_Object  
87  
88 End_Object  
89
```

AI Agent (Inactive)

`Customer.Purchases` directly as a `cWebColumn` with a currency mask.
Let me now make the edit:
Hmm, the column was inserted but it ended up after `oColStatus`. Let me fix the ordering so **Total Revenue** sits before **Active**, and also clean up the extra blank line:
Let me also clean up that extra blank line between the two columns:
Let's do a final read to confirm the whole file looks perfect:
Everything looks great! Here's a summary of what was added:

New column — `oColTotalRevenue` (line 73–79)

```
Object oColTotalRevenue is a cWebColumn  
  Entry_Item Customer.Purchases  
  Set psCaption to "Total Revenue"  
  Set piWidth to 130  
  Set pbPromptButton to False  
  Set peAlign to alignRight  
End_Object
```

Detail	Explanation
Field	<code>Customer.Purchases</code> — the running total of all order amounts for each customer, maintained automatically by <code>cOrderHeaderDataDictionary</code>
Position	Between Phone and Active — a natural place before the boolean column
Alignment	<code>alignRight</code> — standard for currency/numeric values
Currency formatting	Inherited from the DD, which sets <code>MASK_CURRENCY_WINDOW</code> on <code>Purchases</code> via <code>cCustomerDataDictionary</code>

claude-sonnet-4-6 · 127,408 in · 2,435 out · \$0.4187

Great. Can you try to compile it?

Sure!

Compile current project?
Compiling saves all unsaved changes before running the normal Studio compile command.
Compile the project to verify the new FlexRoadsDemo.wo view with the Total Revenue column builds without errors.

Waiting for response...

DATAFLEX STUDIO AI AGENT

THE DATAFLEX-SPECIFIC TOOLS

- `read_code` — read source files in the open workspace (from the editor buffer).
- `edit_code` — insert/modify lines in the Studio editor.
- `compile` — run the DataFlex compiler.
- `read_output` — read errors and warnings.

DATAFLEX STUDIO AI AGENT

THE DATAFLEX-SPECIFIC TOOLS

- `read_code` — read source files in the open workspace (from the editor buffer).
- `edit_code` — insert/modify lines in the Studio editor.
- `compile` — run the DataFlex compiler.
- `read_output` — read errors and warnings.

- `command_line` — execute terminal commands.
- `search_docs` — search through the documentation for extra context
- `create_breakpoint` — create breakpoint on a line.
- `debugger` — go through breakpoint step by step and read it out.

DATAFLEX STUDIO AI AGENT

WHERE THIS IS GOING

- Today: it can do small, well-scoped tasks reliably.
- Soon: longer, multi-file changes, refactors, test generation.
- Progress is about better tools and better context.

/*END*/

