

Bad Prompts, Bad Plans: Teaching AI to Understand PostgreSQL



Thursday, 25 June and Friday, 26 June 2026

OST Eastern Switzerland University of Applied Sciences, Campus Rapperswil (Switzerland)



Bertrand Hartwig-Peillon


Working at **HUG**
(Hôpital Universitaire de Genève)

Daily playing with 500 Postgres instances (mostly docker based) - we use Postgres from 2019.

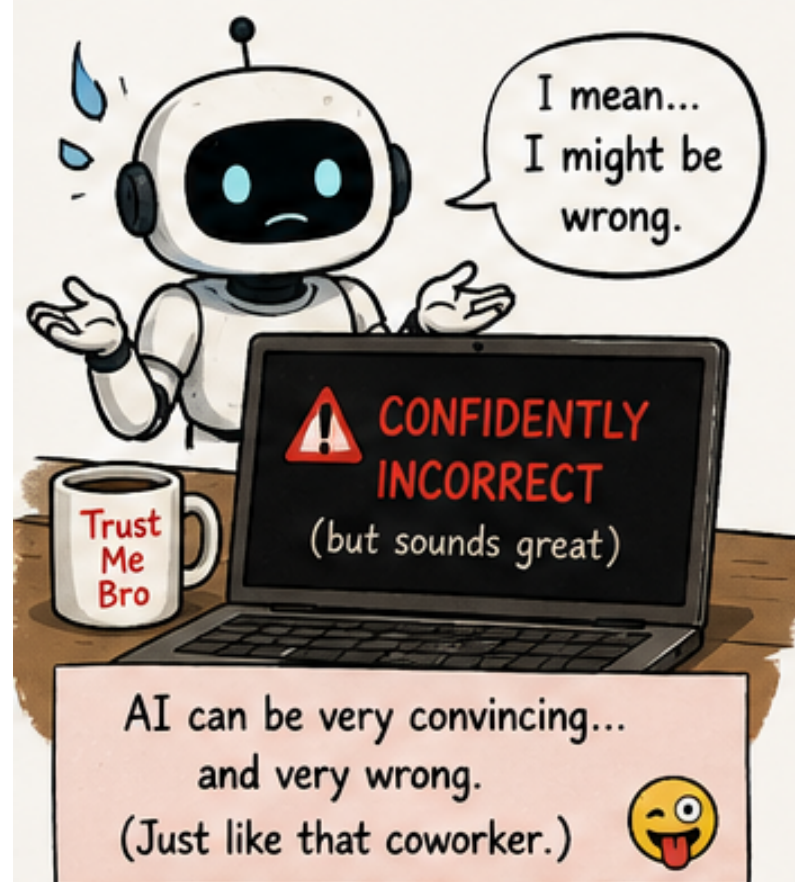
I am working on the OSS tool **pgAssistant** when I find time for it (nights and week ends) and when i'm not fishing or hiking.



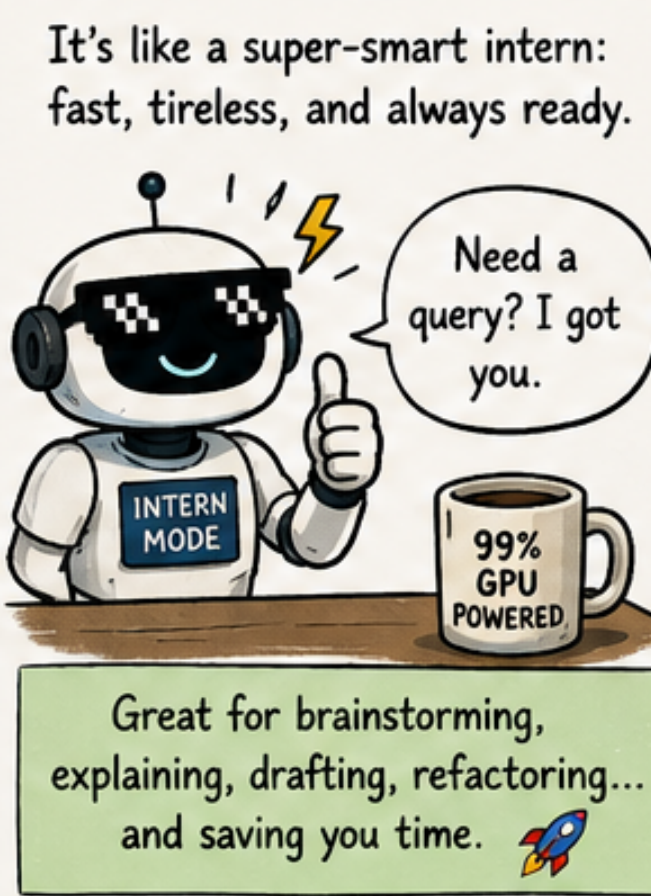
Why this talk ?



❌ **NO**, AI is not a source of truth.



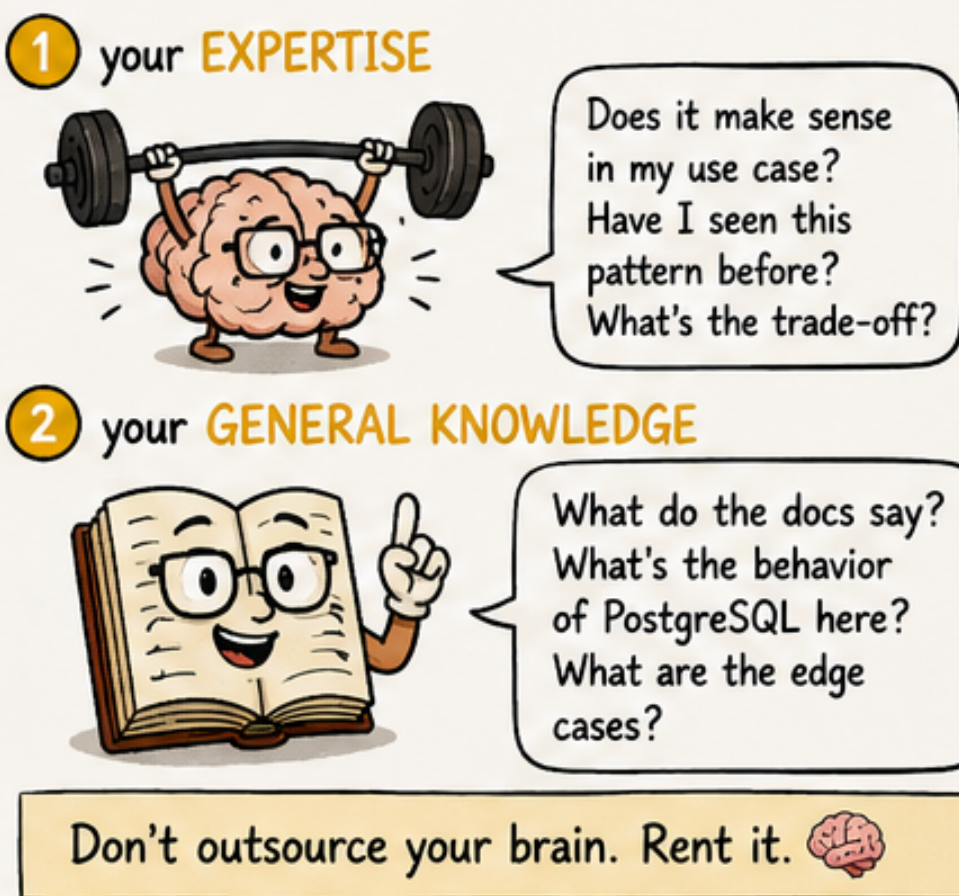
✅ **YES**, it can help you.



❓ But don't ask questions without minimum context.



★ And most importantly, **challenge** the answers with:



AI is your **copilot**, not your **autopilot**.

Context + critical mind = useful answers.

TL;DR:
Think. Verify. Understand. Then execute. 😎

40 years ago, what we learned ...


Remember expert systems 40 years ago.



Philosophers have long observed something fundamental about humans:

Humans are cognitively lazy.

And in practice, that observation is true.



1. THE FIRST TIME

An expert system tells you: "Go left."



The first time, you question the recommendation.

2. AFTER A FEW CORRECT ANSWERS


After a few correct answers, you start trusting the system.



After a few correct answers, you start trusting the system.

3. BY THE TENTH RECOMMENDATION

By the tenth recommendation, you no longer even read it carefully — you simply execute it.

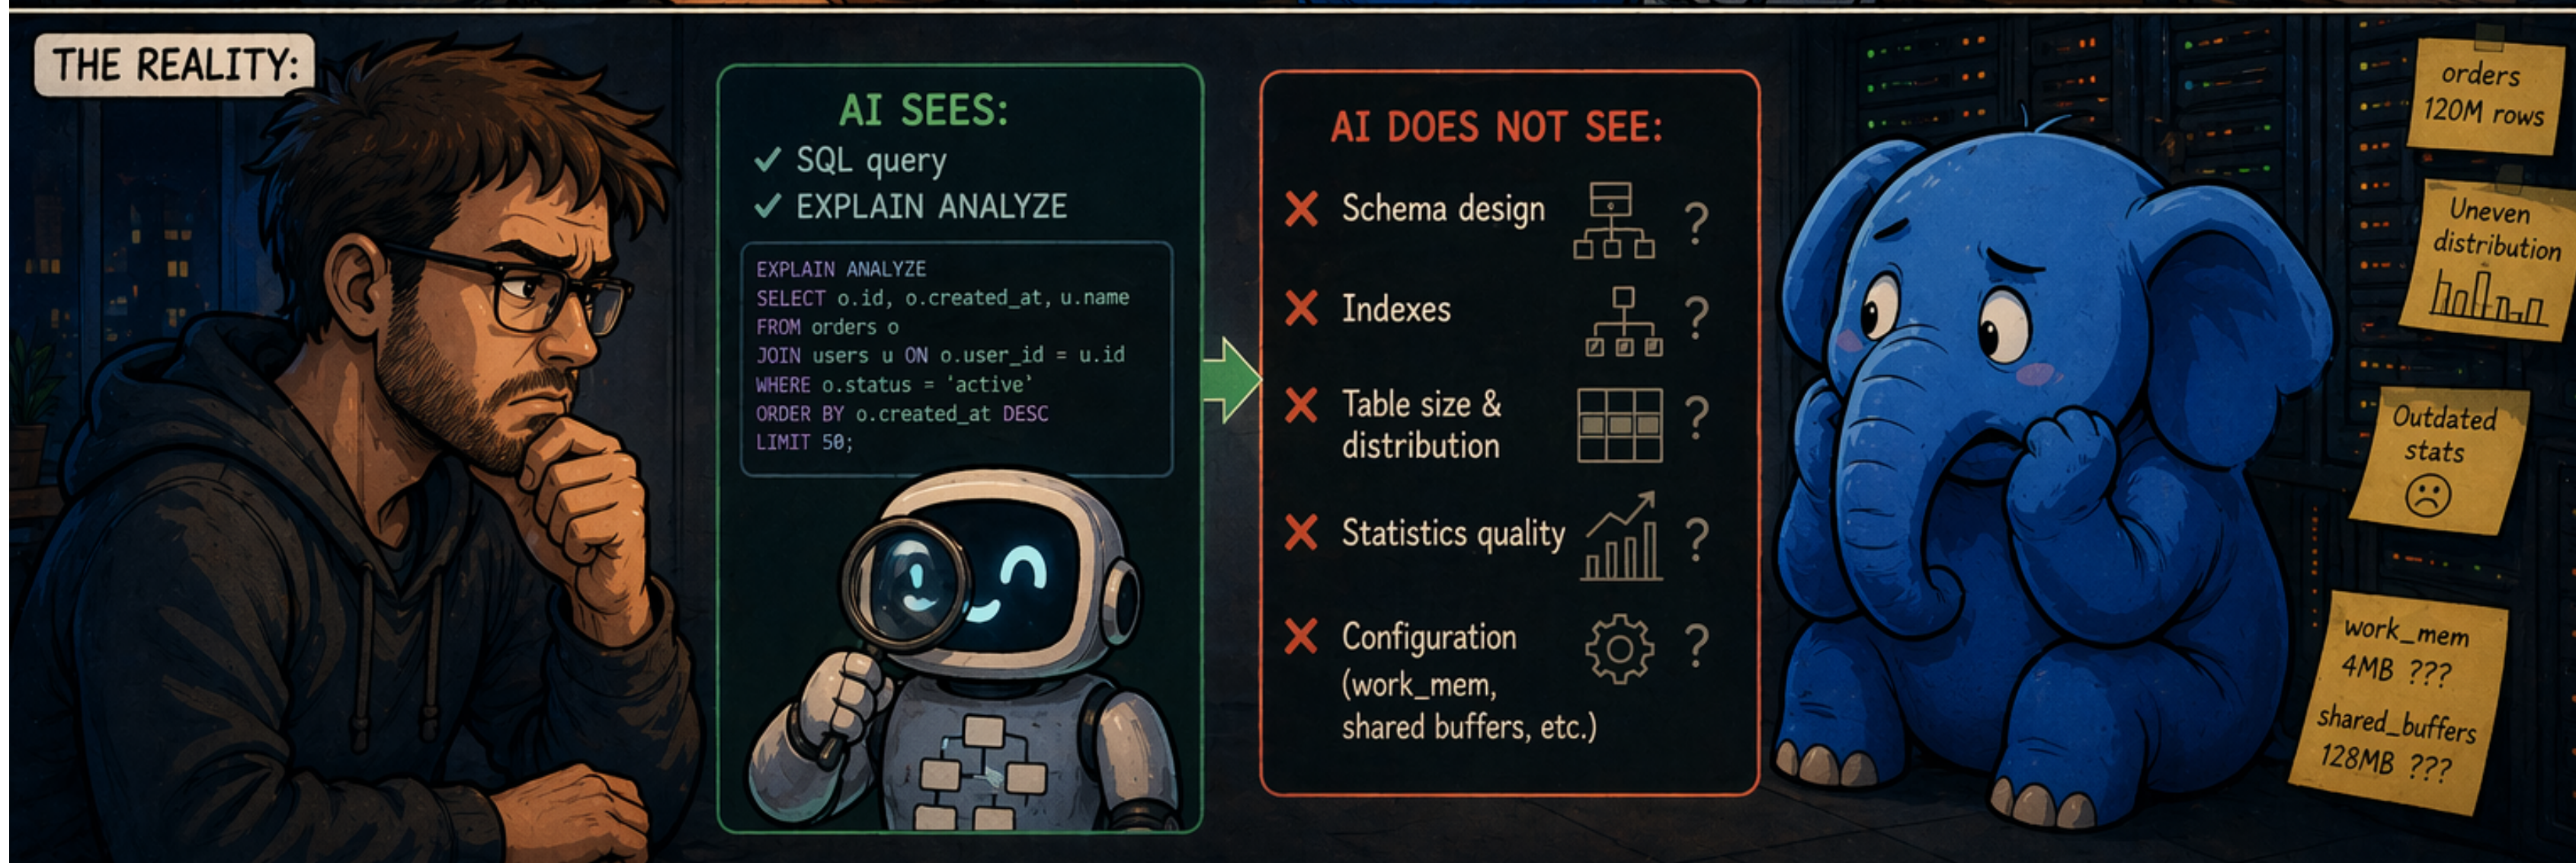
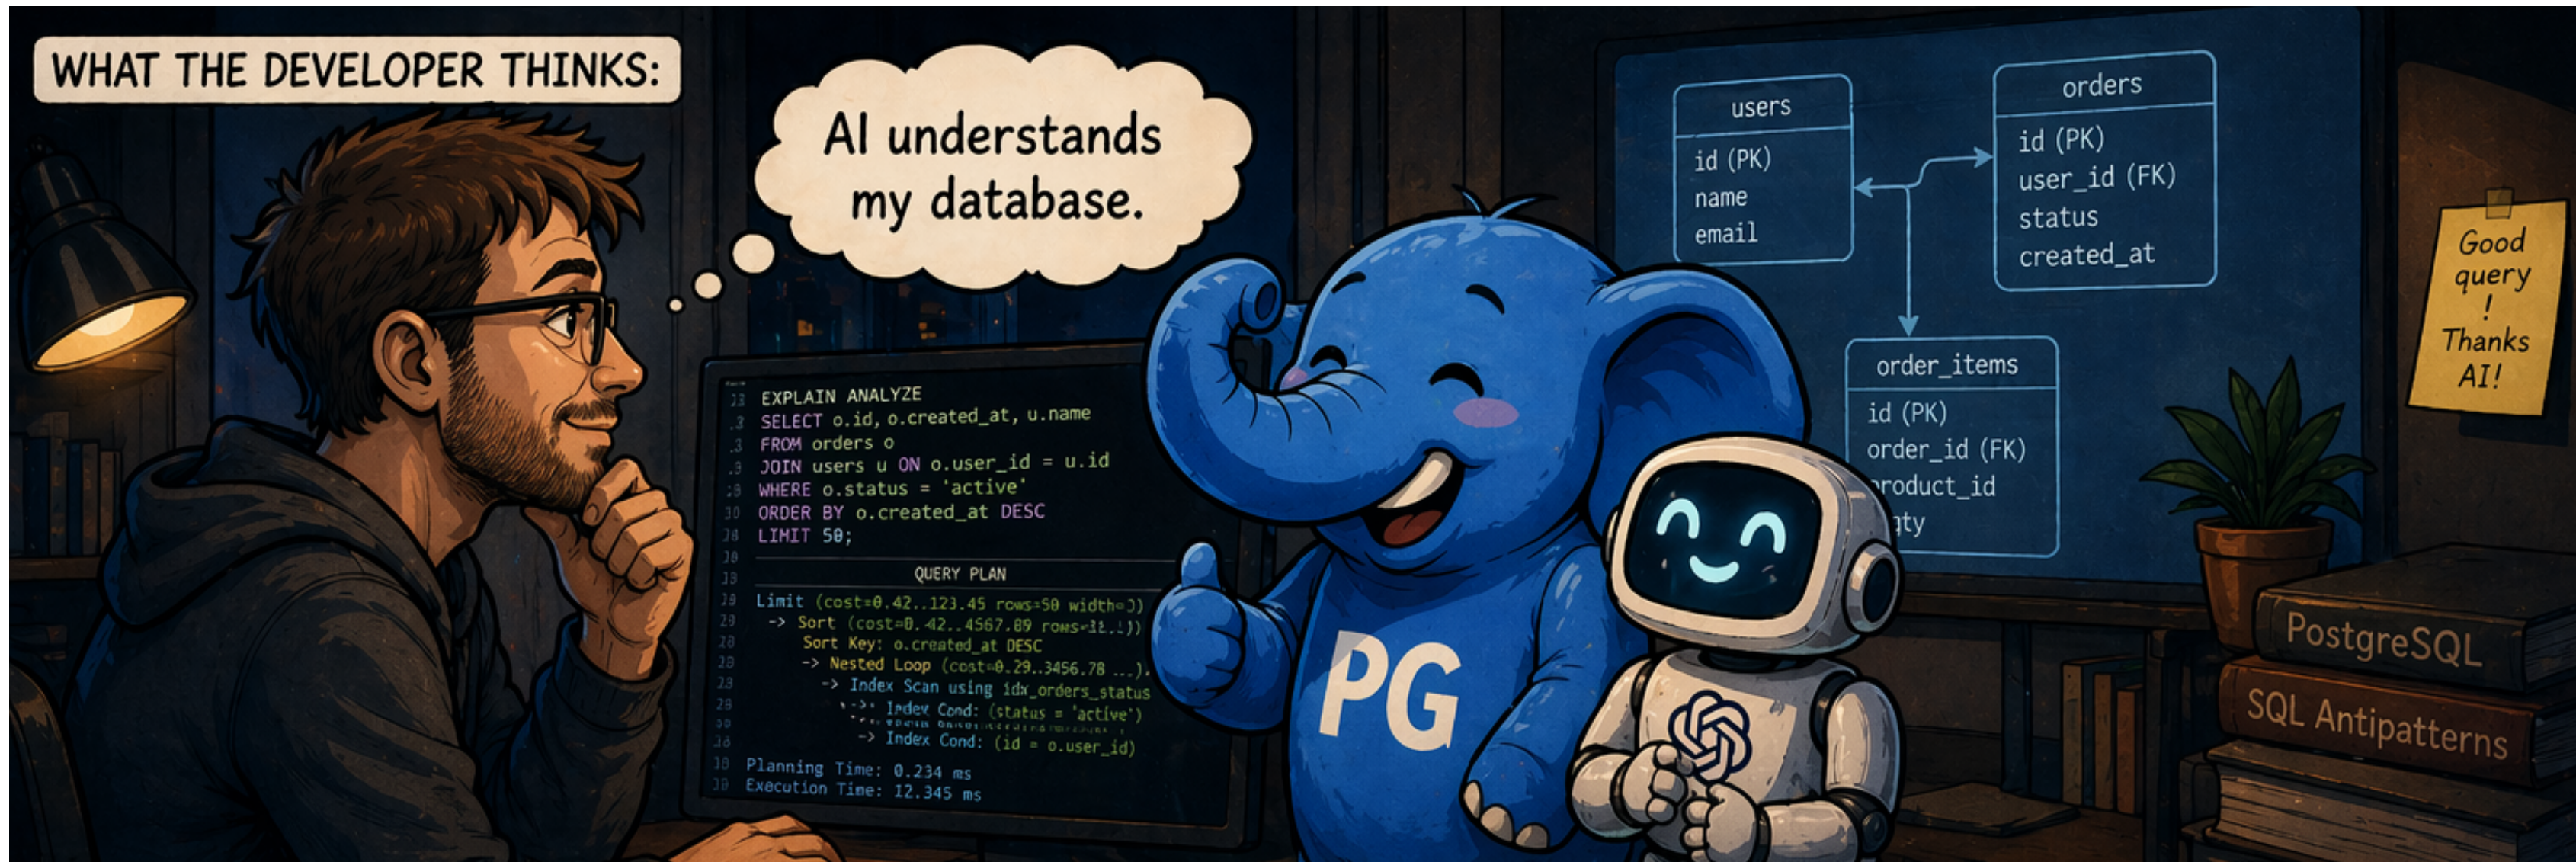


By the tenth recommendation, you no longer even read it carefully — you simply execute it.



THAT IS WHERE THE DANGER BEGINS.

Dev mindset in 2026



DBA mindset



pgAssistant and AI

THE PROBLEM WE JUST SAW...

AI understands my database.

AI ONLY SEES THIS:

- SQL query
- EXPLAIN ANALYZE

AI DOES NOT SEE:

- ✗ Schema design
- ✗ Indexes
- ✗ Table size & distribution
- ✗ Statistics quality
- ✗ Configuration (work_mem, shared_buffers, etc.)

THE SOLUTION: GIVE AI THE FULL CONTEXT

With **pgAssistant**, I provide the missing context:

- ✓ SQL query
- ✓ EXPLAIN ANALYZE
- ✓ Schema design (tables, columns, relations)
- ✓ Indexes (all details)
- ✓ Table size & distribution
- ✓ Statistics quality
- ✓ Configuration (work_mem, shared_buffers, etc.)
- ✓ PostgreSQL version & settings

Now AI has everything it needs to help me!

Developer + **pgAssistant** → Complete prompt → Smarter AI

THE RESULT: BETTER ANSWERS, REAL IMPACT

AI RESPONSE (WITH FULL CONTEXT) ✓

Root cause:
The query performs a nested loop on a large table because the filter is not selective enough and there is a missing composite index.

Recommendation:

```
CREATE INDEX idx_orders_user_status_created_at ON orders (user_id, status, created_at);
```

Why:
This will allow the planner to use an index scan and avoid scanning 120M rows.

Additional suggestions:

- Increase statistics target on orders.status
- Consider adjusting work_mem for this workload
- Refresh table statistics (ANALYZE)
- ...

Better context,
Better answers,
Better performance! ❤️

pgAssistant

The open-source tool that builds complete, context-rich prompts for AI to analyze your PostgreSQL queries.

- 📄 All the context AI needs
- 🎯 Better recommendations
- 🔧 Smarter performance tuning
- 🔗 Open source & developer friendly

Less guesswork
✓ More accuracy

Real impact
✓ on performance

Built for PostgreSQL lovers ❤️


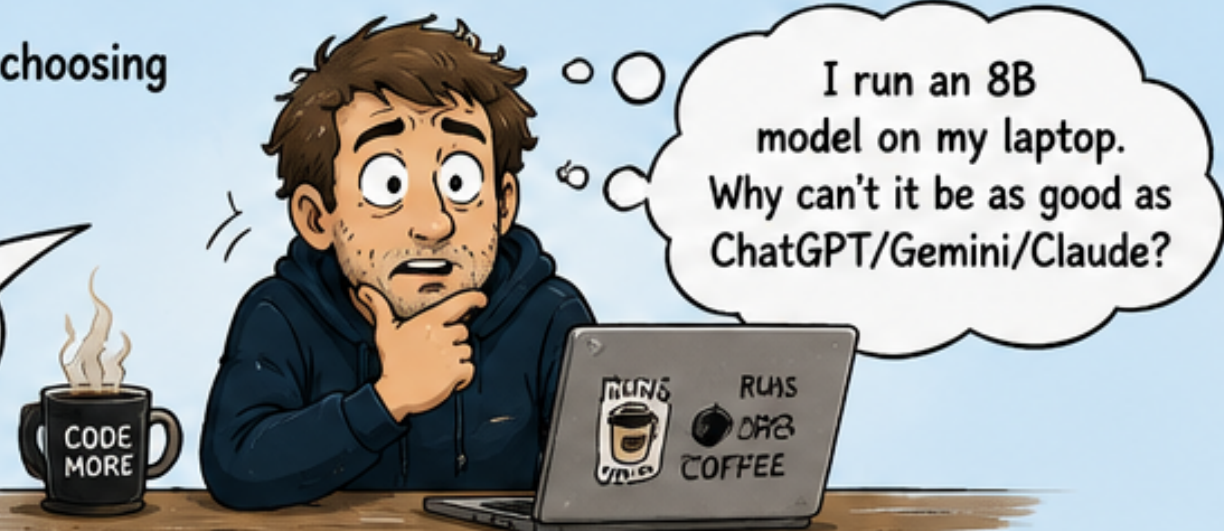
github.com/beh74/pgassistant-community

The Right AI for the Right Query

NOT ALL AI ARE CREATED EQUAL

Choosing the right AI is like choosing the right car for the job.

DON'T EXPECT A SCOOTER TO WIN A FORMULA 1 RACE!



SOME AI ARE JUST NOT BUILT FOR THE JOB


It's not just about size. It's also about specialty.



Great with poetry. Not great with query plans.



Great with hype. Not great with facts. *Results may vary. A lot.



Great with code. Not great with Postgres.



Very confident. Very often wrong.

8B ON YOUR LAPTOP

The scooter



- ✗ Fits in your backpack
- ✗ Cheap to run
- ✗ Great for simple tasks
- ✗ But... has limits

ASK: Explain PostgreSQL execution plans in detail.

ANSWER: Sure! Basically, it works like... ummm... joins are like... tables together? 🤔

20B MODEL (CLOUD)

The sporty car



- ✓ Much smarter
- ✓ Better reasoning
- ✓ Handles complex tasks well
- ✓ Great all-rounder

ASK: Explain PostgreSQL execution plans in detail.

ANSWER: Here's a detailed explanation of the plan nodes, costs, cardinality estimates, indexes, and what could be improved... 😎

120B+ MODEL (CLOUD)


The Formula 1



- ✓ Deep reasoning
- ✓ Highly accurate
- ✓ Understands nuance
- ✓ Handles edge cases like a pro

ASK: Explain PostgreSQL execution plans in detail.

ANSWER: Here's a comprehensive analysis of the execution plan, including node-by-node breakdown, potential issues, advanced tuning strategies, and trade-offs... 😍



THE RIGHT AI FOR THE RIGHT JOB

- ✓ Choose the right model size for your needs and budget.
- ✓ Use specialized AI for specialized tasks.
- ✓ Bigger isn't always better, but tiny is often not enough.
- ✓ Context, data, and prompt quality still matter (a lot!).



Pick wisely, ask smartly, and always verify.

AI is your assistant, not your oracle.



The Secret Sauce Behind pgAssistant

1 TO GIVE THE LLM THE DDL OF THE TABLES INVOLVED IN THE QUERY, pgAssistant USES TABLES THAT ARE GIVEN BY THE QUERY PLAN, NOT BY THE QUERY ITSELF



The query uses a **VIEW**, but the plan shows it actually touches **4 tables!**

WAAPIE: THEN SQL QUER? (uses a view)

```
SELECT c.name, SUM(s.amount) AS total_spent
FROM sales_by_customer sbc
JOIN customers c ON c.id = sbc.customer_id
WHERE sbc.total_spent > 1000
ORDER BY total_spent DESC
LIMIT 10;
...
```

EXPLAIN (ANALYZE, BUFFERS)

```
...
HashAggregate (cost=...)
  Group Key: c.id, c.name
  -> Hash Join
    Hash Cond: (oi.order_id = o.id)
    -> Hash Join
      Hash Cond: (o.customer_id = c.id)
      -> Seq Scan on customers c
      -> Hash
        -> Seq Scan on orders o
        -> Hash
          -> Seq Scan on order_items oi
          -> Hash
            -> Seq Scan on products p
            ...
```

TABLES USED (extracted from the PLAN)


	customer	c	-- customer master data
	orders	o	-- customer orders
	order_items	oi	-- items in each order
	products	p	-- product catalog

By using the plan, pgAssistant collects the DDL for these 4 tables to give the LLM the full context.

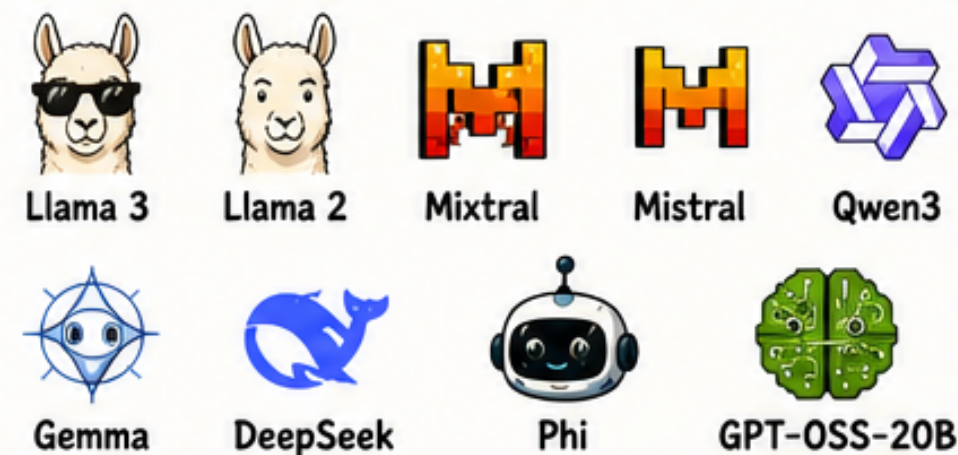
2 USING A LARGE LANGUAGE MODEL (LLM) EFFECTIVELY REQUIRES MODEL-SPECIFIC CONFIGURATION.

Parameters such as temperature, context window, and system instructions can significantly impact the quality and reliability of the output, and these settings often need to be tuned differently depending on the model being used.

Different models, different knobs! I tune them for you.



pgAssistant auto-configures these models for you:



Key parameters we adjust (per model):

TEMPERATURE

Controls randomness. Too high = creative (but risky) Too low = boring (but safe)

CONTEXT WINDOW

How much the model can "see". Every model has its limits.

SYSTEM INSTRUCTIONS

Define the role, style and rules. A small change, a big impact!

Different models, different sweet spots. pgAssistant handles it for you.


3 ADDITIONALLY, PROMPT LENGTH PLAYS A CRITICAL ROLE:

overly short prompts may lack necessary context, while excessively long prompts can dilute important signals or exceed the model's optimal processing capacity. Striking the right balance in prompt size is essential to achieve accurate and consistent results.



Not too short. Not too long. **Just right!**

TOO SHORT



Huh? Not enough info...

Missing context = risky or incorrect answers

JUST RIGHT



Perfect amount of context. Let's go!

All the right info, clear signal, better answers

TOO LONG



Too much noise... I'm lost!

Diluted signal or exceeds optimal capacity

The goal: enough context to be useful, not so much that it hurts.


QUALITY > QUANTITY



pgAssistant gives the LLM the right context, uses the right settings, and keeps prompts just right. You get better answers, faster.



Demo



What i learned coding pgAssistant ?

Using a expert system (determinist) is a first good approach;

Using AI (not determinist) is a **complementary** approach;

But when an expert system gives you a false recommandation, it's a bug. When LLM gives you a false recommendation, it's a feature.

And in both cases, you still need to challenge the system responses!

pgAssistant, since the beginning is using this 2 approaches and follows this philosophy.

What are the real advantages of AI with PostgreSQL? Two areas stand out:

- 1/ Checking compliance with RFCs,
- 2/ Checking naming conventions.

In both cases, it's difficult, if not impossible, to adopt an "expert system" approach.

What about query optimisation ? it depends : scooter or formula 1 and ... the context length ... If the context length is large, today, a formula 1 can allucinate. Tomorrow ? i don't know.

Keeping historical datasets (Global advisor and query ranking) is not useful, it is mandatory. —> pgassistant-collector and Grafana.



What about the future ?

Autonomous PostgreSQL system ??

The future is more likely to be hybrid ???

A practical architecture for autonomous PostgreSQL might include:

- an LLM to understand a developer's or DBA's request expressed in natural language;
- PostgreSQL monitoring tools to collect objective metrics such as slow queries, locks, wait events, cache hit ratios, replication lag, index usage, etc ;
- rule-based checks to enforce known best practices around schema design, indexing, configuration, backup policies, security, and high availability;
- statistical or machine-learning models to detect anomalies, forecast capacity needs, or identify unusual workload patterns;
- a knowledge base plus a machine-learning containing internal runbooks, past incidents, architecture decisions, and environment-specific constraints;
- a human PostgreSQL expert to validate risky recommendations, especially for production changes.

This is often more robust than relying on a single large model that is expected to diagnose and optimize everything on its own like ... an **Oracle AI** ?



A Closing Thought

I once had a conversation about AI with a renowned professor at HUG. He told me :

AI will make good doctors
excellent — and bad doctors even
worse.