AI performs poorly when context is lacking: Causes and solutions.

The gap between what AI promises and what it actually delivers today is significant. The same AI model can produce accurate and useful results in one system, while in another it may yield generic, irrelevant, or almost useless responses.

Where does the problem lie? The cause is often not the AI ??model. What makes AI perform poorly is the 'context'.

Much of today's enterprise system is not designed for AI to operate. Data is scattered across various tools, user recognition systems lack consistency, and behavioral signals arrive slowly or sometimes not at all. Platforms only record events, but cannot connect them into a continuous stream that fully reflects user behavior.

Meanwhile, AI relies heavily on that continuity. When context is lacking, the model will 'fill in the gaps' with speculative inferences. The final result may look polished, but in reality, it lacks relevance and doesn't accurately reflect the user's actual needs. This is where most businesses currently get stuck when deploying AI.

AI cannot fix poor-quality data.

Many people believe that simply using a more powerful AI model can solve the problem. But the article emphasizes the opposite: a better model cannot fix fragmented, outdated, or disconnected data.

According to Gartner's estimates, organizations are losing an average of about $12.9 million annually due to poor data quality. AI doesn't solve this problem—it only makes it more apparent and on a larger scale.

The author describes AI as a magnifying glass. If the data system is strong, AI will make it much stronger. But if the data is weak, fragmented, and inconsistent, AI will expose all those weaknesses. Organizations that once 'lived with' fragmented customer data can no longer hide the problem with slow reporting or manual analysis.

'Mirror test': the quickest way to test AI systems

There's an interesting testing method called 'mirror testing'. The idea is simple: give the AI ??a clear customer intent signal and see how it responds. If the output remains generic or irrelevant, the model may need improvement. But if the AI ??works perfectly with clean data, then completely fails when fed into real-world production data, the problem lies with the data, not the AI. In practice, this almost always falls into the latter category.

Even if the data quality problem is solved, a bigger change is still underway: the way customer profiles are built is completely transforming.

Previously, enterprise data systems primarily stored 'static content'. CRMs stored transactions, data warehouses stored demographics, and marketing platforms stored campaign responses. This data only described what had happened in the past. It was suitable for reporting, but not designed for AI.

AI needs context — and context isn't static data. The context is the current picture of the customer, including:

1. Recent behavior
2. Multichannel signal
3. The intention is taking shape.
4. The chain of connections between successive interactions

For example, if you ask the AI ??to suggest a beach vacation spot, it might suggest Hawaii or Florida. But if the AI ??knows you have three young children, the results will shift to be more family-friendly. And if the AI ??also has access to your recent search history, financial indicators, and search trends over the past year, the suggestions will be completely different because the model is no longer based on general demographics, but is looking at a 'living picture' of you as you are right now.

That's the difference between identity and context. Identity tells the AI ??who you are. But context tells the AI ?? what you're doing and what you're likely to do next.

The real problem lies in the system architecture.

The biggest challenge lies not in theory but in the technical infrastructure. Context doesn't exist within a single system. It's fragmented across event streams, product analytics, CRM, data warehouses, and various real-time pipelines. For AI to effectively utilize context, businesses must shift from batch processing to a streaming or near-real-time architecture.

In other words, user signals need to be collected, processed, and delivered to the AI ??almost instantly.

This is also why many AI projects stall. The model is ready, but the 'context layer' isn't properly implemented. The system can't retrieve the appropriate signals within milliseconds or resolve identity across multiple channels in real time. Until that happens, context remains a theoretical concept rather than something AI can actually use.

MCP is helping AI maintain continuous context.

Architectures like MCP (Model Context Protocol) are also helping to accelerate this process. MCP allows AI systems to transmit user 'memories' across multiple applications, creating a continuous context chain revolving around each individual through various interactions.

As a result, user profiles will become increasingly richer and more predictive over time, because AI not only knows what a person has done in the past, but also understands what they are doing now and what they might do next.

When the identity and context layers are strong enough, the same AI model will produce much better results. But if these layers are weak, almost no model will be able to compensate.

Another noteworthy point is that AI models are gradually becoming 'interchangeable'. The real competitive advantage will belong to organizations capable of operating context at scale and viewing the model as merely a 'processing layer', rather than the entire advantage.

Businesses that invest early in identity systems, first-party data, real-time data infrastructure, and context synchronization mechanisms are beginning to reap cumulative advantages that are difficult to catch up with. Good data helps train better models. Good models attract more users willing to share data. And the more quality data, the stronger the system becomes over time.

This gap is structural, not simply algorithmic.

The future of AI will revolve around context.

The organizations that win in the AI ??era won't be those that write the best prompts. Instead, they will be those companies with systems good enough to understand customers even before the prompt is written.

This means that the future direction of AI investment will also change. Instead of just chasing new models, businesses will have to focus more on:

1. Real-time data
2. Identity author
3. Streaming architecture
4. Context retrieving
5. Reliable consent management system

These are less flashy than a new model launch, but they are far harder to replicate — and may well be the real advantage AI will have in the coming years.