An Enterspeed white paper

10 things about controlling AI and their hallucinations

By CEO Toke Lund

www.enterspeed.com

Contents

Intro

The question I hear most often these days isn't "Should we adopt AI?"
It's "How do we trust it?" And rightfully so.

I’ve yet to meet a retailer using large language models that hasn’t encountered hallucinated outputs. A product description that invents a specification. A campaign text that contradicts your brand guidelines. A customer-facing copy that gets the price wrong. These aren’t edge cases – they’re a systemic challenge that every enterprise scaling AI must confront.

Here’s the thing most people miss: AI hallucinations aren’t bugs in the traditional sense. They’re a feature of how language models work – predicting the most statistically likely next word rather than the most factually correct one. When you understand this, you realise that controlling hallucinations isn’t about fixing AI. It’s about architecting the environment in which AI operates.

That’s what this paper is about. Ten things that actually work when you’re deploying AI at enterprise scale – written for the people responsible for making it work in practice, not just in principle.

Enterspeed CEO, Toke Lund

1. Context is everything – but not all context is equal

The single most powerful hallucination mitigation strategy sounds obvious: give AI access to the right data at the right time. This is where most enterprises stumble. They either flood models with irrelevant information (creating noise) or starve them of necessary context (forcing the model to guess).

For a retailer, this plays out every day. Your AI is generating product descriptions, but it’s pulling from a data model that hasn’t been updated since last season. Or it has access to everything – including deprecated SKUs, old supplier copy, and conflicting specifications – and doesn’t know which source to trust.

> “...give AI access to the right data at the right time.”

The solution isn't more data. It's smarter data architecture. Think of it as creating a data supply chain designed specifically for AI consumption.
Your AI needs structured, validated, and contextually rich information that's dynamically served based on the query. When models pull from authoritative sources in real-time, hallucination rates drop dramatically.

The companies getting this right are treating their data infrastructure as strategic AI assets, not just storage systems. They're building what you might call "AI-ready data" – information that's tagged, transformed, and made immediately accessible through standardised protocols.

For retail specifically, this means your product data, pricing, inventory, and content all flowing from a single, trusted source.

2. Human-in-the-loop isn't realistic at scale (but human-on-the-loop is)

You can’t approve every AI output. But you can build systems that tell you which ones need your attention.

Let's be honest: you can't have a human approve every AI output when you're processing thousands of product descriptions, campaign texts, or customer interactions daily. It’s just not doable. But abandoning oversight entirely isn’t the answer either.

The right model is “human-on-the-loop” – intelligent flagging systems that identify which outputs need human attention and let the rest flow through. Think of it like quality control in manufacturing. You don’t inspect every unit. You build processes that catch anomalies and flag edge cases.

In practice, this means building confidence scoring mechanisms that understand when a model is uncertain, when it’s operating outside its training domain, or when it’s dealing with high-stakes outputs – like compliance-sensitive product claims or legally required information.

There’s a catch, though… you can only act as a human-on-the-loop if your data layer carries the necessary metadata to flag risks automatically.

Without proper data architecture, you’re back to manual review or flying blind. The companies getting this right have invested in tagging, structuring, and enriching their data so the system knows what’s sensitive, what’s verified, and what requires a human eye.

3. Your training data is only half the battle

A model trained six months ago doesn't know about this season's range, your new pricing structure, or last week's regulatory update.

Everyone obsesses over training data quality, and they should. But production AI systems face a different challenge: staying current.

This is particularly acute in retail, where product ranges, pricing, promotions, and compliance requirements change constantly. Your beautifully trained model from six months ago doesn't know about this quarter's product launches, regulatory changes, or organisational updates.

This is where retrieval-augmented generation (RAG) approaches shine. Instead of relying purely on what the model "remembers" from training, you dynamically inject current, relevant information at the moment of generation. The result: AI that reflects your actual business reality, not a snapshot from several months ago.

There's one important caveat that I should mention: RAG is only as good as your retrieval system.

If you're pulling irrelevant documents or outdated information, you've traded one problem for another. The retrieval mechanism needs to be intelligent, fast, and deeply integrated with your business logic – which brings us back to data architecture as the foundation.

"Your beautifully trained model from six months ago doesn't know about this quarter's product launches, pricing changes, or new compliance requirements."

4. Data governance isn't just about compliance – it's hallucination prevention

Poor governance doesn’t just create legal risk. It creates AI that contradicts your own business reality.

Imagine your AI assistant serving up product information for a line you’ve discontinued. Or generating content based on a supplier spec sheet that was superseded six months ago. Or worse – using pricing data from a market you no longer serve.

These aren’t just compliance problems. They’re examples of what you might call organisational hallucinations, where the AI’s output conflicts with your actual business reality.

Proper data governance creates boundaries that prevent AI from accessing or surfacing information it shouldn't. This means implementing real-time policy controls that understand data sensitivity, user permissions, and regulatory requirements.

When your AI infrastructure has these guardrails built in at the data layer, you’re not just preventing compliance violations. You’re preventing a whole class of hallucinations where AI makes up responses because it lacks proper context about what it can and can’t reference.

The enterprises doing this well have moved beyond static access controls to dynamic, context-aware data policies.

5. Standardisation is your friend (even when it's boring)

The dirty secret of enterprise AI is that a lot of teams are rebuilding the same things. Custom integrations for every data source. Bespoke connectors for every tool. One-off implementations that are impossible to scale. This fragmentation doesn’t just waste engineering time – it creates hallucination vectors, because every integration is a potential point of failure.

Without standardisation, you're building legacy debt with every new AI integration. Each custom connector becomes technical debt that someone has to maintain, update, and eventually replace. Worse, these one-off solutions make it nearly impossible to implement consistent validation and governance across your AI systems.

Open standards like the Model Context Protocol are changing this dynamic. Instead of writing custom code to connect your AI to Google Drive, then Slack, then your database, then your CRM, you build against a single protocol. Think USB-C for AI: one standard that works everywhere.

This standardisation matters for hallucination control because it allows you to implement consistent validation, monitoring, and governance across all your data sources.

When you have uniform interfaces, you can build uniform safety measures. Plus, the time you save on integration work can be redirected to what actually matters: prompt engineering, testing, and validation.

“Without standardisation, you're building legacy debt with every new AI integration.“

“When you have uniform interfaces, you can build uniform safety measures. That’s the real value of standardisation.”

6. Not all hallucinations are created equal

A content brainstorm tool and a product compliance tool should not have the same guardrails.

There's a crucial distinction between creative hallucinations and factual hallucinations. In some contexts like brain-storming, campaign ideation, generating product naming options a degree of speculation is actually useful.

The problem arises when you need factual accuracy, but your system is configured for creativity.

This is why temperature settings, system prompts, and model selection matter so much. A customer service bot should be configured conservatively, but a marketing ideation tool can run hotter.

Smart enterprises build different AI configurations for different use cases, each with appropriate constraints. They don’t deploy one-size-fits-all AI. They build task-specific implementations with controls matched to the risk profile.

See table on the next page.

Use case type	Example use cases	Hallucination tolerance	Typical controls	Business risk
Creative / Exploratory	Ideation, brainstorming, marketing concepts	High	Loose prompts, high temperature	Low
Informational / Operational	Internal assistants, support responses	Medium	RAG, confidence scoring, source validation	Medium
Critical / Regulated	Legal, financial, medical, compliance	Near zero	Human approval, audit logs, strict governance	High

For retail, product compliance content, legal text, and pricing information sit firmly in the “critical” category.

Campaign ideation and draft copy sit in the “creative” category. Treating them the same way is where things go wrong.

7. Monitoring and feedback loops are non-negotiable

If you're not monitoring your AI outputs systematically, you're flying blind.

Monitoring AI outputs means more than logging requests and responses. It means tracking hallucination patterns, measuring output quality drift over time, and identifying which types of queries consistently produce unreliable results.

The best systems implement real-time telemetry that detects when model behaviour starts diverging from expected patterns – before it reaches customers or gets published. This requires building feedback mechanisms where editors, marketers, and operations teams can flag incorrect outputs, which then feed into improvement cycles.

You wouldn’t deploy a new ecommerce feature and never monitor its performance. AI systems need the same discipline – arguably more, because AI behaviour can drift in ways that traditional software doesn’t.

8. Explainability isn't just nice to have – it's a control mechanism

When you understand why AI made a decision, you can catch errors before they compound.

When you can trace why an AI produced a particular output – which data sources it weighted, which context it prioritised – you can spot hallucinations faster.

This becomes particularly powerful when combined with domain expertise. A product manager who can see which parts of the data the model weighted heavily can quickly identify when it's over-indexing on irrelevant information.

In retail, explainability matters at every level: compliance teams need audit trails, editors need to understand why a description was generated a certain way, and leadership needs confidence that the system is operating as intended.

Dimension	Black box AI	Glass box AI
Visibility	None	High
Data sources	Unknown	Explicit and traceable
Decision logic	Opaque	Interpretable
Error detection	Reactive	Proactive
Human oversight	Difficult	Built-in
Regulatory readiness	Low	High
Trust level	Assumed	Verifiable

Image by freepik

9. The temperature setting everyone forgets

Temperature isn't a technical detail. It's a business decision about acceptable risk.

Model temperature controls the creativity-accuracy trade-off. Higher temperatures make models more creative – and more likely to hallucinate. Lower temperatures make them more deterministic, which is exactly what you want when generating product specifications, compliance copy, or legally required text.

The mistake most teams make is using default settings across all use cases. Your product content tool should probably run at a low temperature setting. Your campaign ideation tool can run much higher. These aren't just technical parameters – they're decisions about how much creative latitude your AI is allowed to take with your brand.

The key is matching temperature settings to your specific use case and then testing rigorously. Find the sweet spot between useful creativity and reliable accuracy for each application, and document it as part of your AI governance framework.

10. Trust, but verify – at the architectural level

The enterprises succeeding with AI aren't the ones who've eliminated hallucinations.
They're the ones who catch them before they cause damage.

The fundamental principle underlying all of this: no matter how good your controls are, you need verification mechanisms baked into your architecture. At every layer.

The retailers moving fastest with AI right now aren't the ones who've found a way to trust AI blindly. They're the ones who've built layered systems that catch and correct errors before they reach a product page, a campaign, or a customer. They architect for failure – because they know failure will happen.

Trust is built not by assuming correctness, but by verifying it at every level.

On the next page, you'll see what that looks like when you translate it from principle into practice.

What this means in practice

- ✓ Automated validation against authoritative data sources for factual claims
- ✓ Confidence scoring that surfaces uncertain outputs for review
- ✓ Human approval workflows for high-stakes, customer-facing content
- ✓ Audit trails so you can trace exactly how an output was generated
- ✓ Staged rollouts where new AI capabilities are tested at limited scale before full deployment

Layer	**Purpose**	**Hallucination risk addressed**
Human oversight	Final accountability	High-stakes errors
Monitoring	Detect drift & patterns	Silent degradation
Model configuration	Control creativity	Over-speculation
Context retrieval	Ground answers	Knowledge gaps
Governance & access	Enforce boundaries	Unauthorised data use
Data quality	Trust the inputs	Garbage-in hallucinations

Figure: Defense-in-depth for AI hallucination control.

Hallucination mitigation cannot rely on a single control.

Effective systems apply layered defenses across data quality, governance, context injection, model configuration, monitoring, and human oversight.

Each layer reduces risk, and together they create resilient, auditable AI architectures.

So, what should you do next?

The good news: you don’t need to reinvent the wheel. The tools, the standards, and the architecture patterns exist. What’s needed is the will to implement them systematically – and the discipline to treat data quality and governance as strategic priorities, not IT hygiene.

The companies getting this right are discovering that the same infrastructure that reduces hallucinations also makes their AI faster, more accurate, and more valuable. Better data architecture isn’t just about risk mitigation. It’s about unlocking what AI can actually do for your business.

Trustworthy AI isn’t about having perfect models. It’s about having the right systems in place for when they’re imperfect. And they will always be imperfect.

Start with three questions:

- Where is your AI currently generating outputs without a reliable, real-time data source to ground it?
- Which of your AI use cases carry the highest business risk if they hallucinate – and do they have the right guardrails?
- Do you have monitoring in place that would tell you if output quality started to degrade?

If you’re not sure where to start, we’d be happy to walk you through how other retailers are approaching this – including how Enterspeed’s data orchestration layer sits at the foundation of a reliable AI content supply chain.

About Enterspeed

Enterspeed accelerates innovation by minimising complexity.

We bridge new and legacy systems, and give enterprises a single, trusted data supply chain for AI, digital orchestation, and agentic commerce.

Trusted by enterprises across Denmark, Norway, and the UK.

Our mission? Accelerate growth. Reduce complexity.

We enrich and structure your data delivery
and break down silos to ensure a single source of truth for GenAI and digital operations.

We decouple and abstract systems
eliminating dependencies on rigid architectures and providing full flexibility.

We optimise for GenAI
by delivering lean, structured, and real-time data to any model.

We reduce costs and complexity
ensuring lower operational costs by minimising maintenance, infrastructure overhead, and development effort.

And above all... **we empower digital teams**
with a simple, API-driven approach that enhances speed, scalability, and ease of integration.

Enterspeed

Accelerate growth. Reduce complexity.

Get in touch with us anytime
info@enterspeed.com
+ 45 21 55 18 25

Enterspeed A/S, P.O. Pedersens Vej 2, DK-8200 Aarhus N, Denmark