

Saving Automotive with Agentic AI

Current state of opportunities, challenges, and recommendations

Jan Zawadzki | CEO @ Stealth (Reliable Agentic Workflows)

10/07/26

I've worked for 10+ years in AI and help executives understand the nooks and crannies of AI to achieve ROI.


- Served as Head of Artificial Intelligence at Cariad SE, Volkswagen AG's central software development company, and CTO of CertifAI GmbH, focused on AI testing
- Established and led AI engineering hubs in Germany and China, focusing on safe and scalable AI solutions for infotainment and autonomous driving
- Completed MD assignment recently at AHEAD Automotive, a Joint Venture between Clarios, Hella, ZF, and Niterra. Has 8+ years of experience in the automotive industry and 10+ years working in AI.
- Created the [AI Project Canvas](#) and [AI Solution Canvas](#), a globally used and publicized tool for AI project management
- Member of multiple Data & AI leadership committees and guest lecturer at Oxford University and ESMT Berlin
- [Widely-read author](#) on "Towards Data Science" platform and [Keynote Speaker](#), with articles and public appearances garnering hundreds of thousands of views
- Former mentor and content reviewer for [Deeplearning.AI](#). MBA from ESMT Berlin and MSc in Computer Science.



Current State of AI


AI is incredibly powerful, yet
something doesn't add up

AI shows signs of “jagged intelligence”, surpassing humans on narrow benchmark, but failing on broader benchmarks.



While AI is superhuman on some tasks, it’s unreliable in broader categories. Safety improvements on model side lead to degrading accuracy.

AI Investment and Adoption rates are sky-high, but ROI is absent.



While money is spent on AI and it is being adopted, PnL impact isn't realized. Reliability is the core issue.

Please scan the code below and complete the survey yourself.

<https://form.menti.com/s/alkanwz1ofds>



Making Agentic AI Work

Management Illusions About
Agentic AI


Agentic AI can be this centuries' most powerful technology.

Agentic AI is the next logic development going from general learning algorithms to autonomously acting software.



AI: it predicts

- **What it is:** The umbrella term for software that learns from data to classify, predict and recommend.
- **Examples:** Spam filters, card-fraud alerts, Netflix recommendations and CRM lead scoring.
- **Superpowers:** Finds patterns in millions of data points and decides in milliseconds.



Generative AI: it creates

- **What it is:** AI trained on vast amounts of text, images and code that creates new content on request.
- **Examples:** ChatGPT, Claude or Gemini drafting proposals, summarizing calls and writing code.
- **Superpowers:** Turns a plain request into a first draft in seconds, in almost any language.



Agentic AI: it acts

- **What it is:** GenAI that plans multi-step tasks and uses tools to finish them with little human input.
- **Examples:** Coding agents that fix bugs end to end; sales agents that research leads and book meetings.
- **Superpowers:** Works toward a goal end to end, across your systems, around the clock.

We've discovered ten management illusions about making AI work for Enterprises and Hidden Champions.

- Across industries, companies are experimenting with agentic AI, yet no one has really figured out how to make it work.
- This paper was written after discussing around this topic for 6+ months with the Heads of AI from Ergo, EnBW, 50 Hertz, Miele, appliedAI, CARIAD, and CLAAS, stemming from the "Data & AI Leaders Cycle" in Germany.
- The results of the discussion led to the creation of the set of management illusions that we've encountered. We've grouped the illusions by stage.
- We then validated the illusions through an online survey with technical talent.
- The paper is submitted for peer-review and expected to be published in October.




Fig 1: Management Illusions along the Agentic Implementation Journey

Beware of the 10 Management Illusions when selecting and scaling AI-based and Agentic AI Solutions.

AGENT-FIRST ILLUSION

Agents should be integrated into every process and product.

END-TO-END ILLUSION

Agents can complete at sufficient quality

TESTING ILLUSION

Agents work reliably out-of-the-box, testing is an afterthought.

SCALING ILLUSION

If an agent performs well in one use-case, we can easily scale it across use-cases or the company.

DOWNSIZING ILLUSION

We can reduce headcount if agents take over tasks.

PRODUCTIVITY ILLUSION

If we automate a process, we become more productive overall.

BOTTOM-UP ILLUSION

AI champions from each department will figure out how to make agentic work.

AGENT=EMPLOYEE ILLUSION

We need to treat agents like employees and include them in org charts.

MASTER PLAN ILLUSION


We need an "Agentic Transformation", akin to a "Digital Transformation".

PREPARATION ILLUSION

We must either just get started with Agentic AI or develop the perfect plan.

AI-first E2E Business Processes Cause Similar Challenges.

At each step of business automation, errors can happen, compound, and lead to brand damage



TESTING ILLUSION

AI that doesn't behave as
intended loses companies
money.

Integrating AI into your product means integrating a loose screw - it will break. How do we tighten it?

The most underdiscussed, most important challenge of making AI work

- As an engineer or business owner, you can't blindly trust the output of the AI-based products.
- A lack of reliability leads to unintended behavior, with unaccounted for consequences.
- How do we evaluate and put guardrails on AI-based applications, to use them in safety-critical and also business-critical operations?

How do we tighten the loose screw of AI?



Source: Shutterstock/Jiranart Aonnoum, accessed Jun 5th 24

Revenue is dependend on a great customer experience.

Testing is indispensable for developing a high-quality product and for offering a superior customer experience.

TESTING TRTHOUGHOUT THE
WHOLE PRODUCT LIFECYCLE



1



Testing throughout the **whole product lifecycle** based on **highest state-of-the-art industry standards**

TESTING

2



The product meets all the **relevant benchmarks** as **proven by testing** resulting in a **high quality**

QUALITY

3




The high-quality product **maximizes customer experience**

CUSTOMER
EXPERIENCE

4-Tiered Approach to Reliable AI Engineering

Structured testing of the AI application is the main prerequisite for a reliable AI engineering approach.



The concept of edge case testing on the boundary of an application space is common in software testing.


Defining the Application Scope

- The application scope defines the boundary of the application
- Similar to SW testing, to test the AI application properly, we need to test its behaviour on the edges of its application scope.
- Robustness is indicated by achieving the desired performance on the boundary of the application
- Carefully defining the test set based on the application scope is paramount to test for robustness

Example: Credit Worthiness Score

- To test the performance of on AI model that predicts the credit worthiness of applicants between the age of 55 – 65, and with a salary between €80k – €100k, we can define the ODD as shown on the right
- We can then place test objects strategically in the test set to test for robustness and performance on its edges

Defining the application scope is model- and data-agnostic and should be implemented as best practice.



The management illusions have different levels of severity. Managers should focus on the top-right set of illusions.

Validating the results with technical experts

- The results of the survey of 167 technical experts confirm that our illusions exist in their eyes.
- Managers should focus on not falling for the downsizing and testing illusion.
- The agent=employee illusion seems to be the least prevalent and negatively impactful illusion.




Fig 2: Results of the management illusions study, ranked by prevalence and negative impact of illusion


Way Forward

How to make AI Work for
You

AI is neither Automotive's savior, nor henchman. Be nuanced around making AI work for you pragmatically.

How does Automotive win through Agentic AI?

- Be pragmatic about implementing Agentic AI.
- Start with your business strategy.
- Beware of the 10 Management Illusions about implementing Agentic AI.
- If you can solve them without AI – do that! If you need AI, don't hesitate.
- Accept that AI will never be perfect.
- When someone shows you an AI PoC, always ask: Did you do a **risk analysis**? **How did you test it**? Do we have domain-specific data?
- Be smarter than your competition about using AI, use it coming from a **good engineering** background.
- **Not using AI is not an option.**



Contact

For more information regarding Agentic Blue or Agentic Reliability, please contact:

Jan Zawadzki

jan.m.zawadzki@pm.me

+49 176 3152 6527

This document has been prepared for discussion purposes only. All information is subject to and qualified in its entirety by the terms of a definitive agreement and a signed Non-Disclosure Agreement.

Selecting the right use-cases to implement is the most important step to guarantee AI ROI.

1

The Agent-First Illusion

- Description: AI agents should eventually enter virtually any business process or product.
- Impact: Having a hammer and treating every problem as a nail. Agentic AI brings unpredictability and variance into process and product quality that needs to be managed. If it can be avoided, don't use agentic AI.
- Example: After one agentic AI use-case ideation workshop, only 10% of the proposed ideas should be solved with agents, the rest could be solved with deterministic workflows.

2

The End-to-End Illusion

- Description: AI agents can complete complex tasks from beginning to end automatically.
- Impact: Suboptimal performance of the use-case and very hard way to realize ROI on an initiative.
- Example: Trying to use an agent for the entire sales process. If the agent makes an error at one-step, this error propagates through the process. Rather, use agents for specific tasks where needed, and embed them in more deterministic software.

If an agentic workflow works in the first demo, this doesn't mean it can work as a product or that it scales.

3

The Testing Illusion

- Description: AI agents work reliably and out-of-the-box. Code agents create the right tests.
- Impact: Use-cases repeatedly fail in production. Engineers firefight to locate issues and fix them. Quality of the product is subpar, and adoption stalls, and users churn.
- Example: A voice agent works well when tested in a specific setup with a call center. When calling customers, the agent failed to perform as expected.

4

The Scaling Illusion

- Description: Managers underestimate the challenges of scaling one agentic workflow to other use-cases, or even company-wide.
- Impact: Management assumes that with limited investment, the workflow can be scaled across the company and use-cases. Disappointment arises when this fails.
- Example: An agentic workflow exceeds expectations. It was selected, because it was working with the cleanest underlying data. Scaling this workflow to messier data sources failed.

Just because tasks can be automated with AI, it doesn't mean that jobs can be eliminated or that productivity rises.

5

The Downsizing Illusion

- Description: Managers inflated expectations that staff can be reduced when a share of their tasks can be automated.
- Impact: Managers plan for staff layoffs and when they cannot happen, reorganization structures and financial planning is affected.
- Example: Companies like Klarna or Salesforce prematurely lay off customer support and other staff, only to find that quality of the replacements is subpar.

6

The Productivity Illusion

- Description: AI agents work reliably and out-of-the-box. Code agents create the right tests.
- Impact: Making certain isolated processes in a company more efficient doesn't mean the whole company becomes more efficient. The bottleneck might simply shift to the next not automated process.
- Example: An invoice processing agent classifies invoices and saves many hours of work, yet finance staff still has to verify the outcomes and create the financial planning manually.

We recommend creating a top-down initiative and to not anthropomorphize AI.

7

The Bottom-Up Illusion

- Description: Managers belief that a grassroots movement of “AI champions” will push agentic AI forward across the company.
- Impact: Companies generate many ground-up ideas, but departments lack the budget, competence, and legal support to develop them. Lack of governance and support from the top stalls enablement on the lowest levels.
- Example: One company with 30,000 employees dedicated 60 FTE from the CIO department to AI transformation, aligned closely with HR, and managed to perform the transformation efficiently.

8

The Agent = Employee Illusion

- Description: Managers urge to anthropomorphize agents and to integrate them into the organizational chart.
- Impact: Agents aren’t used as a tool, but the processes around managing humans are applied to managing agents without a clear benefit.
- Example: A company created personalities for their agents to mirror the corresponding employees. The company even created management layers for their agents with feedback cycles. The results was token burn without improved performance.

Noone has figured out the perfect playbook for agentic AI yet. Be pragmatic about your strategy.

9

The Master Plan Illusion

- Description: Managers belief that a master plan is needed to get started, akin to “digital transformation” an “agentic transformation”.
- Impact: Managers start to create KPI’s around deploying agents, instead of looking at the differences between departments and treating it as a tool.
- Example: One company set out the goal to have an agent for every employee. Another company created the goal to deploy 10,000 agents working directly with their ERP system until the end of the year.

10

The Preparation Illusion

- Description: Managers belief that you can either start implementing agentic AI without a strategy or they need a perfect infrastructure before getting started.
- Impact: Companies either run into a wall when deploying agents or don’t get started at all. Both approaches lead to frustration and wasted budget.
- Example: One company developed an AI strategy for over six months before the advent of agentic AI. It then had to be reworked, without having gained any experience working with AI.