

AI IN THE MITTELSTAND

AI in the German Mittelstand: How Strong Ideas Shape the Future

Artificial intelligence is considered the key to the productivity of German small and medium-sized enterprises (the Mittelstand). But what can language models and agents really deliver - and where are the limits? A conversation with Dr. Damir Dobrić, co-founder, CEO and Lead Software Architect of daenet GmbH - ACP Digital, about realistic use cases, the future of software development and missed opportunities in the digitalization of government.

Handelsblatt (sponsored content) – <https://www.handelsblatt.com/adv/firmen/ki-im-mittelstand.html> – English translation of the German original

About: Dr. Damir Dobrić

Position Co-founder, CEO and Chief Architect of daenet GmbH

Company daenet GmbH (founded 1999 in Frankfurt am Main, today part of the ACP Group)

Other roles Professor at Frankfurt University of Applied Sciences, Microsoft Regional Director and Microsoft Most Valuable Professional (MVP)

Focus areas AI and cloud architecture, digital transformation and modern software development

Expertise One of the most experienced and recognized AI and cloud architects worldwide

Dr. Dobrić, hardly any topic occupies boards and managing directors as much as artificial intelligence right now. Where does the German Mittelstand really stand today - beyond the headlines?

Honestly, in a phase of experimentation. Almost every company has by now tested a chatbot or let a few employees experiment with ChatGPT. But there is a world of difference between an impressive demo moment and productive, economically measurable operations. The Mittelstand recognizes the potential - but often does not get beyond the first experiments.

Why is that? Usually there is a lack of clear use cases, of cleansed and usable data, and of secure integration into existing systems. Companies do not need another AI demo. They need a realistic roadmap from pilot project to productive operation. Companies also lack the expertise. A course of a few weeks at the employment agency does not turn anyone into an expert.

Let's stay with the basics. What can large language models - the technology behind ChatGPT and the like - really deliver today?

Quite a lot already, if you know what for. Language models are excellent at understanding and generating language: they summarize documents, answer questions about large bodies of knowledge, translate, structure unstructured information and draft texts. Wherever language, knowledge and patterns are involved, they are enormously productive.

That is no small thing. A large part of classic office work consists of exactly such activities - and this part will largely be taken over by AI in the coming years. Tools like Microsoft 365 Copilot already show this in everyday work - in writing, summarizing and researching.

And what can they not do? Where should decision-makers become skeptical?

A language model does not understand in the human sense. It calculates the probability of which word comes next. That means: it can sound convincing and still be wrong. A pure language model is not built for genuine logical reasoning, for guaranteed correct figures or for reliable facts.

It is important to understand what language models cannot do. They can neither directly control a refrigerator or a vacuum cleaner, nor play chess. They are merely able to generate text and to derive from the user's intent which action (tool) should be executed.

The actual action, however, is not performed by the model itself, but through an interface provided by developers, which controls the corresponding system and executes the desired function.

For this reason, the supposed responsibility of an AI is also frequently misunderstood – especially in political discussions and comparable bodies. The language model does not make independent decisions about executing actions; it merely provides an interpretation of the user input. The responsibility for the actual execution and its effects lies with the implemented software (tool) and the developers who provide these interfaces.

That is why my most important advice to all decision-makers is: build up technological competence.

Every new technology requires a certain degree of understanding. Only those who know its possibilities and limits can make well-founded decisions. If this understanding is missing, the probability of wrong decisions is high.

A negative example of this are cookie banners. They impressively demonstrate the consequences of regulatory decisions taken without sufficiently understanding the underlying technology. The result is a solution that makes daily life harder for millions of people without effectively solving the actual problem.

One term comes up more and more often: “AI agents”. What does that mean – and why is it more than just another chatbot?

A chatbot usually follows a simple request-response pattern: the user asks a question, the model answers, and the interaction ends until the next input. An AI agent, by contrast, works in a continuous loop. Thanks to its so-called reasoning capability, it can analyze a task and break it down into simpler subtasks. It then performs various actions: for example, it queries a database, sends an email or calls another system. It evaluates the results and repeats the process until the goal is reached or it terminates the operation.

However, agents only unfold their full value as part of an ecosystem – with access to company data, business applications and tools, as well as security and governance mechanisms. So that a model can extend its untrained knowledge and connect to external systems, so-called MCP servers are used: they provide functions and interfaces (tools) that merge with the AI. This is exactly where agent platforms come in. Microsoft Foundry offers a platform for developing, deploying and managing AI agents; with the Microsoft Agent Framework, their logic and multi-step workflows can be implemented and orchestrated. Such platforms ensure security, control, model independence and maintainability, and provide the functions in a way that most companies are already familiar with.

You rely heavily on Microsoft - Azure, Foundry, the Agent Framework. Why this clear commitment?

Because in enterprise use, governance, security and reliability are decisive. And that is exactly where Microsoft plays to its strengths. Most medium-sized companies have long been working with Microsoft 365 and Azure anyway. Docking AI exactly where data, identities and permissions are already cleanly managed is the shortest path to a productive and secure solution. Via Microsoft Foundry we access a large selection of commercial and open-source models while remaining model-agnostic. Add to that topics that are often underestimated in the Mittelstand: data residency in the EU, compliance, responsible-AI guardrails and traceability. A platform like Azure brings all of that with it, instead of you laboriously building it yourself. And with Copilot, AI has long since arrived in the everyday work of many employees: in the document, in the email, in Teams. We at daenet have been deeply rooted in the Microsoft ecosystem for years - I myself as Microsoft Regional Director and MVP. As part of the ACP Group, we bring this platform expertise into every project.

You yourself come from software development. How much are agents like Claude or GitHub Copilot changing the way software is created?

A great deal. I have been developing software for more than 30 years, and I have rarely experienced an upheaval at this pace. These tools write code, explain unfamiliar code, find bugs, propose entire architectures. What used to take days is now created in hours. The speed of implementation has increased enormously.

But - and this is the decisive point - speed of implementation is not the same as speed to a finished, reliable product.

For example, this summer we are launching a completely new process platform called "Wingman", which first uses AI intensively to understand processes from video and speech as well as all available information, but then analyzes and improves the process together with humans (agent). Afterwards, this improved process is handed to people for training or into process automation. As a side effect, in the best case, I have secured the company's entire knowledge. We have medium-sized customers with thousands of business processes.

By the way, for a fraction of the cost of asking an LLM for a cooking recipe. That is real value.

You have to explain that. At first it sounds like a contradiction.

It is the most underestimated point in the entire debate. AI helps to implement very quickly. But in doing so it also introduces errors - and often it is not able to fix its own errors. As a human, you have to look into the generated code to correct exactly the problems that the AI itself has caused. You can ask a human who implemented code; you cannot ask an AI.

That means: some of the time you gain at the front end you spend again at the back end - on reviews, refactoring and corrections. These efforts increase significantly and should not be underestimated. They often lead to considerable time losses. In the end, the extremely high efficiency is strongly limited by exactly these losses. In terms of efficiency, the result is simply not as spectacular as many believe.

I am deliberately not saying that AI does not increase productivity - it certainly does. But the industry will still have to learn in the near future that the net effect is smaller than the marketing promises suggest.

Does that mean that in the future anyone will be able to build software - even without developer training?

Partly. And here I have to differentiate clearly. Creating software is becoming significantly more efficient - but most of all for those who are developers themselves. Those who understand what the code does can guide the AI, assess its suggestions and correct its errors. For this group, the productivity leap is the greatest.

Many technically skilled people who are not trained developers will also be able to master one solution or another for themselves - that is valuable, but not new. There are already many so-called low-code and no-code tools that do the same with similar productivity. But as soon as things get stuck - and they get stuck regularly - someone has to look into the code and the process and fix the problem that the AI introduced and cannot solve itself. Without this understanding, you eventually hit a wall.

Those who only bring in an expert at that point usually react too late. An experienced person will prefer to rebuild the solution with the support of AI rather than fix an existing half-finished solution.

That is why AI does not replace developers – it shifts their role from the so-called “code monkey” to skilled, creative engineers and architects.

Let's come to daenet and ACP. You implement AI use cases with your own platform, DeepOcto.AI. What is behind it?

DeepOcto is our modular accelerator – deliberately not a rigid standard product. You can think of it as a building-block solution that at its core uses the Microsoft Agent Framework as open source: the DeepOcto platform contains various building blocks – for data connectivity, for retrieval, for agent logic, for connecting models, for integration into existing systems, for data protection and governance.

Depending on the use case, we combine the appropriate building blocks into a secure, productive solution. That is the decisive advantage: we do not start from scratch with every project, but reassemble proven modules. And because new building blocks are constantly being added, the platform becomes more powerful with every project. This is how we get from concept to working result very quickly.

Do you have a concrete example of what such a use case looks like in practice?

Take an industrial company with decades of knowledge in manuals, maintenance logs and emails – scattered across many systems. One DeepOcto building block connects these data sources and cleanses them; another ensures that a language model reliably draws the right answers from exactly this company knowledge instead of making something up. An agent can then also act – for example, create a ticket or prepare an order.

In doing so, we do not just build an AI agent or copilot, but accompany the entire cycle: data, model selection, agent logic, integration, security. And we proceed in three clear steps – understand, implement, operate and scale – whether in the cloud, on-premises or hybrid. The fact that we not only build solutions but can also operate them reliably across the entire region and let them grow is also thanks to our membership in the ACP Group.

We consistently rely on the Microsoft Azure platform and Microsoft Foundry. There, more than 11,000 AI models are available to us, including models from leading providers as well as open-source models. This allows us to select the appropriate model for each use case and to design our solutions largely

model-agnostic.

Our true competence here is to cast a company's specific requirements into software. We combine AI with creativity and excellent engineering skills.

Your promise is that AI makes software not only faster, but also better. How does that fit with your warning about errors?

That is not a contradiction, but a question of discipline. AI can make software better – it finds weaknesses, suggests tests, documents, checks for patterns that a tired human overlooks. As an additional pair of eyes, it is a real quality gain. On a platform like Azure, security and verification mechanisms come included – from code analysis in GitHub to the governance functions in Foundry.

But “better” only emerges when a competent human takes responsibility for the results. The AI delivers the raw material at high speed; quality emerges in the review. Those who take this seriously get faster and better software.

A side note: I am talking here about software that solves complex problems and automates complex processes. Of course, there is also less complex software, such as a homepage or the like. AI can already create such solutions quite well. Will people also forgo control there? Probably not.

Anyone who believes they can take humans completely out of the loop will get worse software.

What does this mean for employees? There are great fears about the loss of jobs.

I do not believe in appeasement.

Sometimes it is easier to look into the past to predict the future. About 200 years ago, the majority of the population worked in agriculture. The industrial revolution, which began in the late 18th century, and the later motorization fundamentally changed this. The share of people employed in agriculture fell from around 60 percent in 1800 to less than two percent today.

Many old occupations disappeared, but at the same time numerous new and better jobs emerged. After all, someone had to develop and manufacture the tractors, while others sold, operated or repaired them. It will be similar with AI and the technologies that will follow it.

Yes, a large part of office work will change – that is the realistic expectation. Anyone who sorts documents, answers standard inquiries or transfers

information from one system to the next will not be doing this activity in this form for much longer.

The question for every decision-maker is not whether, but how they shape the transition. Relying purely on job cuts is the most expensive path: knowledge is lost that has to be bought back later. The smarter path is to qualify the same people for the new tasks – reviewing and taking responsibility for what the AI delivers. The Mittelstand that takes its people along will win this transformation. New roles and business models around AI will inevitably emerge.

You clearly advocate more AI in the public sector. Where do you see the greatest potential – and what, in your view, is going wrong?

I see enormous potential in the digitalization of public authorities. Hardly anywhere else could so much efficiency and citizen proximity be gained with AI as here. And honestly, I find it a great pity that such tasks are repeatedly awarded to the usual suspects – paid for with taxpayers' money.

It would be more effective to give a larger share of these investments to many smaller software companies rather than only to a few large providers with their sometimes market-dominating framework agreements. They are often faster, more efficient and in some cases more competent than the established large providers and system houses. That would help Germany as a business location enormously. Precisely because we do not do this, we have not made progress on this topic for years. Many fear that the use of AI and automation could cost numerous civil servants their jobs. I see it differently.

Experienced civil servants should use their expertise to develop, introduce and continuously improve digital solutions – with and without AI – until they work reliably and efficiently (see our new process platform Wingman above). They could take on this task until their retirement and thereby make a valuable contribution for future generations.

Civil servants in active service could work much more efficiently with the help of such solutions. As a result, the number of employees in the public sector could be gradually reduced over time. Modern software would support employees in such a way that the administration as a whole works noticeably faster and more effectively than today, with simultaneously decreasing effort and staffing needs.

The natural retirement of many employees in the coming years gives the state the necessary leeway anyway to shape this transformation in a socially acceptable way.

The result would be more capable authorities, faster administrative processes and lower costs for the state.

From a technical perspective, I consider this goal absolutely achievable. I would be ready to take on this challenge at any time. Whether such a project is politically feasible, however, is another matter – I would not put my hand in the fire for that.

That is a clear statement. Isn't it also risky – entrusting small providers with large government projects?

I hear that argument often, and I consider it a pretext. Awarding large contracts to large corporations only creates a feeling of security. The track record of many large public IT projects proves the opposite. And there is another fundamental misunderstanding: specialization does not mean lone fighter. daenet is part of the ACP Group – we combine the focus and speed of a specialist with the substance, reach and operational reliability of a large group. The customer gets both: short paths and a partner who will still be there in ten years. That is security that is real – not just felt.

In the past, the big players were indeed considered the safe choice. Today that is different. Many daenet employees have been with us for more than 20 years. Some of our acquaintances went to companies back then that were considered particularly crisis-proof but have since had to cease operations. The decisive question is therefore: which job ultimately proved to be secure?

Apart from that, there is no absolute security. Neither a state nor a company can guarantee us one hundred percent security. That is an illusion. A permanent employment contract, for example, is often an expression of belief in a security that does not exist in this form.

This may not be directly related to AI. Nevertheless, I believe that AI will fundamentally change and reorder many things. Therein lies a truly great opportunity for all of us.

Fire, too, has brought us many benefits since its discovery, even though we still cannot fully control it today. There is no such thing as one hundred percent security. It is an ideal we should strive for –

while at the same time accepting that we will never fully achieve it.

Risk is managed through clear milestones, through measurable results and through tenders in which smaller, specialized providers also have a fair chance. Technology history shows again and again that small, focused teams can become market-shaping providers: in the consumer market with names like Spotify or Telegram, in the B2B environment with providers like Celonis or DeepL. All started small and all scaled rapidly. That is not a gamble, but the rule.

And there is a second, strategic thought: if the state spreads its investment more broadly, it may help one software company or another to grow into a large enterprise. That is exactly what we lack in Germany. That is an idea worth thinking about – for policymakers.

In closing: What do you advise a decision-maker who wants to get serious about AI tomorrow - without falling for the hype?

First: Do not start with the technology, but with the use case. Look for a concrete, measurable application with a clear quick win – not the prestige project, but what really hurts in everyday operations.

Second: Find a competent advisor or partner who can accompany you through the entire process – from the idea through implementation to operations. That is exactly what we stand for: with the technological depth of a specialist and, as part of the ACP Group, with reliability in operations.

And third: Keep competent people in charge. AI is an extraordinary tool – but it remains a tool. Our guiding principle at daenet is: Think Beyond. We start where others give up. This is exactly the attitude the Mittelstand needs now: curious, competent, willing, courageous and level-headed.