

Research Essay

**Beyond Efficiency: Conditions for Sustainable AI Literacy Training
in German Municipal Administration**

Type of thesis	Research Essay
Course name	DLMAISCTAI01 – Seminar: Current Topics in AI
Course of study	M.Sc. Applied Artificial Intelligence
Author	Korbinian Much
Matriculation number	IU14127772
Tutor	Prof. Dr. Simon Martin
Date	June 2026

List of Abbreviations

AI Act	Artificial Intelligence Act (Regulation (EU) 2024/1689)
DPIA	Data Protection Impact Assessment
DSK	Datenschutzkonferenz (Data Protection Conference)
GDPR	General Data Protection Regulation (Regulation (EU) 2016/679)
LLM	Large Language Model
OZG	Onlinezugangsgesetz (Online Access Act)

Table of Contents

1.	Introduction	1
1.1	Background and Context	1
1.2	Problem Statement	2
1.3	Approach and Method	2
2.	The Challenge: Why AI Adoption Struggles in German Public Administration	3
3.	What Adults Actually Learn: A Theoretical Framework	4
3.1	Adult Learning Is Not School Learning	4
3.2	Sustainable Learning as a Change of Perspective	5
3.3	Competence Develops in Context, Not in Courses	6
3.4	Experience and Self-Efficacy	6
3.5	AI Literacy as Transferable Competence	7
4.	Conditions for Sustainable AI Literacy Training in Public Administration	8
4.1	The Empfinden Workshop as an Illustration	8
4.2	Four Conditions for Sustainable Training	9
4.3	From Conditions to Design	11
5.	Conclusion	12
6.	Declaration on the Use of AI-Supported Tools	14
7.	References	15

1. INTRODUCTION

1.1. Background and Context

The public administration in Germany is often described as being under considerable strain. The reason: It lacks personnel with sufficient technical expertise. In 2024 the German Civil Service Association (dbb Beamtenbund und Tarifunion) stated that 550,000 positions in the public sector were not filled (dbb Beamtenbund und Tarifunion, 2024, p. 5). And this gap will grow larger: By 2030, more than every third public sector employee will retire (Kawik et al., 2022, p. 8) and due to budget cuts most positions will not be replaced. This shortage is not only a problem for big federal ministries. On the local administration level the strains can be witnessed even more clearly. There, small teams have to cover a broad range of tasks and due to the amount of work they lack time, resources and motivation to either look into new tools or introduce them properly at their workplace.

Since July 2024 the revised Online Access Act (German: Onlinezugangsgesetz - OZG 2.0) obliges federal and state authorities to offer their administrative services also digitally. This creates additional pressure to introduce said tools that increase efficiency and the use of artificial intelligence is especially discussed. LLMs can process large volumes of text and generate structured outputs. In theory this makes them well-suited for document-heavy, repetitive tasks in public administration.

A second layer of requirements is introduced by the EU Artificial Intelligence Act (Regulation (EU) 2024/1689), which entered into force in August 2024. Article 4, which applies from February 2025, obliges providers and deployers of AI systems to ensure that personnel develop sufficient AI literacy, understood as the skills and knowledge needed for safe and compliant use (European Parliament and Council of the European Union, 2024, Art. 4). For local administrations, this is not merely a general recommendation but a condition that is likely to shape how such systems can be used in practice.

Taken together, labour shortages, digitalisation requirements, data protection constraints, and new AI literacy obligations point to a more complex problem than is often acknowledged. AI literacy frameworks have largely been developed in educational contexts (Long & Magerko, 2020, p. 2; Ng et al., 2021, p. 2) and have not been systematically applied to the structural constraints of German municipal administration.

1.2. Problem Statement

This paper pursues one central objective: to identify the conditions under which AI literacy training programs lead to sustainable, transferable competence in German municipal administration. We define AI literacy as enabling civil servants to assess, adopt and critically use AI tools independently, both now and later as technology develops. To make this concrete, we draw on the municipality of Empfingen in south-west Germany as a running example throughout this paper. The commune has around 4,400 inhabitants and the Rathaus employs around 10 people. This leads to the paper's central question: What conditions must AI literacy training programs fulfil to allow participants to develop sustainable, transferable competence for working with AI tools?

1.3. Approach and Method

In the following we will discuss the theoretical concept of sustainable AI literacy training in three steps. First, we examine the conditions that currently prevent AI adoption in German municipal administrations. For that we will focus on workforce capacity, legal constraints and fragmented infrastructure. Second, we look at what the adult learning literature says about when training actually leads to lasting competence rather than short-lived tool knowledge. Third, we derive a set of conditions that training programs must fulfil to enable civil servants to independently adopt and use AI tools; not just today, but as the technology continues to develop.

This paper follows a conceptual-analytical approach, meaning it is based on adult learning theory, AI literacy literature and selected work on German public administration. A condition is included when at least two of the reviewed frameworks converge on the same training requirement independently of one another. Based on that we derive design conditions for AI literacy training. The Empfingen workshop is used as an illustration (not as empirical evidence) and shows that the derived conditions are practically possible.

It should also be noted that teaching AI literacy to adults in a professional setting works differently than in a classroom. Civil servants bring years of routine, established workflows and often a healthy scepticism towards new tools. Simply handing them a manual or running a one-day workshop is unlikely to create lasting change (Knowles, 1980, pp. 42–49).

This paper draws on five adult learning frameworks. Knowles (1980) explains why adults learn differently from students in school. Mezirow (1997) asks what actually changes when learning works. Lave and Wenger (1991) show where that learning needs to happen. Kolb (1984) and Bandura (1977) explain how experience turns into transferable skill and why one successful

attempt can change what a person believes they are capable of. Long and Magerko (2020) add the AI literacy perspective: not familiarity with specific tools, but the competence to evaluate and use any AI system on their own.

Together they define what this paper means by effective AI literacy training. They are not separate theories. They build one argument. Each of them will be discussed in more detail in the following sections. This paper does not draw on technology acceptance frameworks. Its focus is on whether civil servants develop the competence to use AI tools independently and critically.

In June 2025 a real AI literacy workshop was conducted for the municipality of Empfingen. Civil servants with no prior technical experience produced working results within a single afternoon. This is not empirical evidence. It illustrates the conditions the framework identifies.

2. THE CHALLENGE: WHY AI ADOPTION STRUGGLES IN GERMAN PUBLIC ADMINISTRATION

Germany continues to lag behind in providing digital public services and developing digital skills (European Commission, 2025, p. 1). This is especially true for AI adoption and AI literacy. The reason for this lies in three major issues: human resources, legal compliance and fragmented digital infrastructure. This paper focuses specifically on small and medium-sized municipalities. The aforementioned challenges are most acute at this level.

Labour shortage makes automation a necessity, but current staff struggle to transform daily tasks into AI-guided workflows. A large share of employees was already working before the release of large language models like ChatGPT. They remain accustomed to paper-based or barely-digital processes. Because of the workload, they lack the capacity to learn the tools that would eventually give them more time. The problem is reinforcing itself and it will not resolve without external intervention.

A second problem is legal compliance. Since July 2024 the revised Online Access Act (OZG 2.0) demands that the digital availability of services is mandatory. The problem: citizens now submit requests digitally, but the work process behind it remains for the most part analogue. The result is a backlog that manual labour can no longer clear. On top of that, the Data Protection Conference (DSK, 2024, p. 6) sets out clear expectations on how citizen data may be processed. Any external AI service that receives municipal data must be assessed for compliance before use. Where the processing involves high risk to individuals, a data protection impact assessment (DPIA) is required under Art. 35 GDPR (European Parliament and Council

of the European Union, 2016, Art. 35) before the service can be deployed. The tools exist, but the legal framework constrains when and how they can be deployed.

Digital infrastructure adds a third problem. Small municipalities cannot simply buy access to commercial AI services. Public procurement law slows that down considerably. State-level IT providers are not yet equipped to offer secure AI environments to all municipalities. IT administrators often see the potential of LLMs but cannot act on it. Mergel (2021, p. 336) describes the result: "the actual implementation is stalling, too complex, not sufficiently coordinated." This applies directly to the three barriers discussed here: workforce shortages prevent capacity for change, legal constraints slow adoption, and fragmented infrastructure makes coordination nearly impossible.

These three barriers explain why AI adoption in German municipal administration is not a question of willingness. Staff often want to use new tools. The problem is structural. Any training program that ignores said barriers will fail. For that it must fit into constrained time budgets, respect legal boundaries and work within the existing infrastructure. A workshop that asks participants to use a commercial cloud-based LLM, for example, fails that test before it starts.

There is also a behavioural dimension. Scepticism towards AI tools is not irrational. It has concrete sources. Work overload leaves no time to experiment. Staff fear losing control over familiar tasks. And many simply do not believe they are capable of working with technical tools (Bandura, 1977, p. 195). Training that ignores these sources of resistance will not overcome them.

3. WHAT ADULTS ACTUALLY LEARN: A THEORETICAL FRAMEWORK

3.1. Adult Learning is Not School Learning

Teaching technical concepts like machine learning principles or neural networks most likely will not provide AI competence to municipal staff. Civil servants in Empfingen process citizen correspondence and calculate rental fees. So their training must start from there and an effective AI literacy workshop needs to connect with the actual problem the employees work on. AI comes only as a layer on top of what they already know. A workshop based on this idea allows the learner to experience the result within the same session. In this paper we derive training conditions from the five frameworks introduced above and apply them to the context of AI literacy in public administration.

One of this paper's main premises is that educating civil servants is different from educating students in school. They have a different sense of autonomy. This difference has an impact on

how AI literacy programs should ideally be designed. Knowles (1980, pp. 42–49) introduced a framework of teaching adults which he called andragogy. In pedagogy we assume a dependent learner who receives new knowledge through a teacher. Andragogy on the other hand assumes an autonomous learner who has prior experience, practical goals in studying and needs to understand the relevance of what is being taught in a new lesson. Knowles also emphasizes that adults need to connect new input to their prior experiences. Otherwise the input will be too abstract and the teaching fails. They will only learn effectively when the input connects to tasks they already perform or need to perform. Knowles et al. (2005, pp. 353–356) later expanded this framework to emphasise that adult learning is most effective when it is centered around a problem and not around specific content, and when students have genuine agency in shaping the learning process.

3.2. Sustainable Learning as a Change of Perspective

A civil servant who attends an AI workshop and learns that LLMs can synthesize documents has acquired new knowledge, but this knowledge will not be enough to change how she approaches her work. To achieve that this new information needs to cause a shift of how she thinks about her own role. If her understanding of her job is defined by specific tasks she performs (scanning a fee schedule, calculating a total) then AI knowledge will not translate into the use of AI. What must change is her conception of the task itself. She needs to understand that her job is not to follow and execute a fixed procedure, but to design a more efficient process. That however is a step which goes beyond simple learning.

Knowles explains how adults learn best. Mezirow (1997) asks what actually changes when learning works. In his theory of transformative learning, Mezirow argues that "new information is only a resource in the adult learning process" (1997, p. 10). For staff in Empfingen, this means more than learning a new tool. It means questioning what their job actually is.

Mezirow calls this process perspective transformation (1978, p. 100). This happens when something which they thought they knew stops working. Mezirow calls this a "disorienting dilemma" (1995, p. 50). When a civil servant who has never written a line of code discovers that she can build a working fee calculator in an afternoon using AI, something shifts, because she had always assumed this was not possible for her. Whether she ever gets that chance depends entirely on how the training is designed.

3.3. Competence Develops in Context, Not in Courses

A workshop that uses a fictional scenario might produce awareness of AI capabilities, but it will not produce true competence in its participants. The goal of a workshop with a lasting impact must be that civil servants will use AI to solve new problems on their own in the future. To achieve that, training must be conducted with real tasks, real data (within the legally permissible boundaries) and with real consequences. The Empfinden case illustrates this: the rental fee calculator was not a training exercise. It was a working tool that staff would actually use. The learning was directly connected to the making.

Lave and Wenger (1991) explain where this kind of learning needs to happen. Their concept of situated learning states that knowledge and competence are not transferable when they are separated from realistic practice. Knowledge and competence need to be co-produced by the learner and the environment in which learning takes place.

Their framework is based on a concept called "legitimate peripheral participation". Newcomers do not learn by being taught, they learn by doing actual work at the edges of a community, taking on more responsibility as they become more capable. What follows from this is less obvious than it sounds: learning, in their account, "is not merely situated in practice, as if it were some independently reifiable process that just happened to be located somewhere; learning is an integral part of generative social practice in the lived-in world" (Lave & Wenger, 1991, p. 35). The context is therefore not just where learning happens. It is, as they put it, "a constitutive element of its content" (Lave & Wenger, 1991, pp. 35–36). A course built around fictional scenarios produces knowledge that works in that course. The same competence developed on a real task, in a real office, under the actual constraints of that office, becomes something a person can actually use.

3.4. Experience and Self-Efficacy

Kolb (1984, p. 21) describes learning as a cycle with four stages: concrete experience, reflective observation, abstract conceptualisation, and active experimentation. The stages do not always follow a clean sequence. But they all need to be covered. A civil servant who builds a working tool and walks away has done something useful. But one who also reflects on why certain prompts worked and others did not, and draws conclusions from that, has built a transferable skill. For that the workshop needs to take the learner through the full cycle Kolb describes (not just the doing-part).

Bandura (1977) adds another aspect: whether participants believe they are capable of doing something at all. Self-efficacy is not a fixed personality trait. It shifts based on experience. The most powerful way to shift it is through what Bandura calls enactive mastery (1977, p. 212), meaning actually succeeding at a task. For civil servants who have long assumed that technology is not for them, a single afternoon of building something that works can be more impactful than any amount of instruction. Not because they were praised or encouraged, but because they did it themselves. This changes what they believe they are capable of. And that is what determines whether new behaviour survives after the training ends.

3.5. AI Literacy as Transferable Competence

Taken together, these frameworks point in the same direction. What matters is not whether participants can operate a particular tool, but whether they have developed the underlying judgement to assess and use any AI system on their own. The AI literacy literature arrives at the same conclusion.

Long and Magerko (2020, p. 2) define AI literacy as "a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace." The emphasis on critical evaluation matters for a practical reason: specific tools change. A civil servant trained to operate one system will need retraining when that system is updated or replaced. What she does not need to relearn is the capacity to ask the right questions of any system, meaning what it can do, where it fails, and when its outputs need checking.

Laupichler et al. (2022, pp. 1, 13) note in their scoping review that the field remains underspecified on exactly this point: what should non-experts actually learn, and how? Carolus et al. (2023, p. 7) offer one answer by breaking AI literacy into measurable components: using AI, understanding its mechanisms, detecting its outputs, and engaging with its ethical dimensions. They also show that these components can be assessed and targeted independently of one another (Carolus et al., 2023, p. 5). For a municipal administrator this is useful. A training program does not need to cover everything. It needs to develop the components that are relevant to the job.

The goal of training design is therefore not familiarity with a set of tools. It is the kind of competence that transfers to new tools, new tasks, and situations that no training program could have anticipated.

These theories are useful design lenses, but they have limits. Knowles' andragogy assumes a degree of learner autonomy that is not always available in hierarchical organisations. Mezirow's transformative learning centres on individual perspective change and says little about what organisational conditions need to be in place for that change to last. Lave and Wenger's situated learning is the most critical case: the concept of legitimate peripheral participation describes gradual integration into a community of practice over time (not a single afternoon). For that reason, this paper does not treat said frameworks as complete theories of organisational change. They are design lenses. Each one points to a condition that makes training more likely to produce transferable competence. Whether that transfer actually holds depends on factors this paper addresses in Section 4, and on the organisational conditions named in the conclusion.

4. CONDITIONS FOR SUSTAINABLE AI LITERACY TRAINING IN PUBLIC ADMINISTRATION

4.1. The Empfingen Workshop as an Illustration

In June 2025, a one-afternoon AI literacy workshop was conducted for staff of the municipality of Empfingen, Baden-Württemberg. Participants had no prior technical background in AI or programming. The task was defined in advance: replace the manual process of calculating rental fees for municipal venues with a browser-based calculator that runs locally, requires no internet connection after creation, and processes no external data.

Using an LLM as a code-generation assistant, participants produced a working HTML/JavaScript application within a single session. The tool reads the complete fee schedule, applies the correct rate based on event type, user group and additional services, and calculates the total automatically. It runs entirely on the employee's own device and transmits no data. The session ended with staff using the tool on a real inquiry. During creation, only the publicly available fee schedule and abstract rule descriptions were entered into the LLM (meaning no booking requests or personal information).

This example is not presented as empirical proof. It shows how the four conditions discussed in the next section can work together in practice: when all four are present, civil servants without technical background can produce genuinely useful AI-assisted tools in a single afternoon. The case also demonstrates a specific mode of AI use (AI as a tool to build tools) that sidesteps the data protection constraints preventing direct use of external AI services on citizen data.

4.2. Four Conditions for Sustainable Training

The theoretical framework in Section 3 points to four conditions that AI literacy training programs must fulfil to produce sustainable, transferable competence in civil servants. Each condition follows from the theories discussed. Knowles (1980) and Lave and Wenger (1991) both require that training is grounded in real work, not abstraction. Mezirow (1997) requires an experience that challenges existing assumptions, which means the task must be genuinely new to the learner. Kolb (1984) and Bandura (1977) require that the learner succeeds at that task and reflects on that success. The legal framework discussed in Section 2 requires that compliance is built into the design, not added afterwards. In the following we will discuss each condition in turn.

4.2.1. Condition 1: Real work context

Training must be grounded in an actual task from the participant's work environment, not a simulated or generalised exercise. Lave and Wenger (1991) are clear on this: the context is not decoration but constitutive of what is learned. A real task means real motivation, real constraints, and competence that develops in a context the learner actually inhabits. For that, effective training begins not with a curriculum but with a problem inventory: which tasks in this office are rule-based, repetitive, and currently done manually?

4.2.2. Condition 2: AI as a tool to build tools, not as a service to consume

Participants must experience AI as an instrument they control, not a black box they query. This distinction is both technical and conceptual. A civil servant who uses a commercial LLM to summarise a document stays dependent: on the service, the connection, and the provider's compliance with data protection requirements. One who uses an LLM to generate code for a locally-running tool produces an artefact they own, can test against known cases, and adapt over time (with appropriate IT support where needed). The latter experience triggers the perspective transformation Mezirow (1997) describes: not "I can use this tool" but "I can build something with AI." This shift is durable in a way that tool proficiency is not. It also directly addresses the data protection constraints that prevent direct use of external AI services on citizen data (DSK, 2024, p. 6; European Parliament and Council of the European Union, 2016, Art. 35). This mode of AI use is not universally the right approach. In contexts without strict data protection constraints, direct LLM use may work fine. In German municipal administration, where citizen data cannot be sent to external services without prior compliance review, building local tools for example can be a more viable path.

4.2.3. Condition 3: A concrete success experience within the training session

Bandura (1977) establishes that self-efficacy is built primarily through enactive mastery: direct personal experience of success. For that, training must be structured so that every participant leaves with something that works. That has real design implications. The task must be scoped to fit the available time. Failure modes must be anticipated. And the trainer's job is not to deliver content but to make sure each participant crosses the finish line. A training session that ends with participants having learned a great deal but produced nothing will not change self-efficacy, and therefore will not change behaviour after the session ends.

A success experience alone is not sufficient. Kolb (1984, p. 21) argues that learning requires reflection after the experience. Without it, the result stays tied to the specific task. A short debrief at the end of the session is enough: which prompts worked, where did the AI fail, and would the same approach work for a different task. That step is what makes the experience transferable.

4.2.4. Condition 4: Data protection and compliance as an integrated design principle

German municipalities operate within a clear legal framework, and no training design can be effective if it requires participants to violate or ignore it. The DSK (2024) guidance makes clear that the use of external AI services involving citizen data requires careful prior assessment. The legal framework is a design parameter, not a barrier. Training programs that begin with said constraint (rather than treating it as an obstacle) will produce more realistic competence. The locally-running tool approach illustrated by the Empfinden case shows how compliance-by-design can be built into the training outcome itself. Participants who understand why data protection shapes the tool architecture are more likely to apply the same reasoning independently when they encounter new AI tools later.

These four conditions are not independent. The conditions reinforce each other. A real task without a success experience leaves participants frustrated. A success experience on a fictitious task does not transfer. Compliance awareness without a working tool remains abstract. By removing one the others lose their effect.

4.3. From Conditions to Design

The four conditions specify what a workshop must make happen. For that, three things need to be decided before the workshop starts: what task, what structure, and what the trainer actually does.

Task selection happens before the workshop. A contact person in the administration identifies one or two tasks that meet the Condition 1 criteria: rule-based, repetitive, currently done manually, and solvable without citizen data as input to the LLM. The rental fee calculator in Empfingen is a clean example. The inputs are a public fee schedule and an abstract event description. No personal data is involved and the output is a number. Tasks that involve individual case judgement, legal discretion, or personal data as direct input do not qualify. Scoping is part of this preparation. Meaning a task that takes an experienced administrator thirty minutes can be turned into a working tool in one afternoon. A task that involves conditional logic across multiple documents probably cannot. The trainer needs to make this assessment in advance. Scope is the main risk factor for Condition 3.

The session opens with a problem, not a lecture. Participants are told: here is the task, here are the constraints, and we are going to build something today. That framing activates the motivation Knowles (1980) describes: they are not there to learn about AI, they are there to solve a problem they already know. They then work with the LLM directly. They describe the task in their own words, inspect the output, test it against known cases, and iterate. The trainer does not demonstrate. He or she moves through the room and makes sure no one is stuck for more than a few minutes. Once the tool works, it is tested against a real inquiry from the office. Not a constructed example. If it gives the right answer, Condition 3 is met. If not, the session continues. A short debrief closes the session. Three questions are enough: what worked, what did not, and where would the same approach apply to a different task. That last question is what makes the experience transferable. Kolb (1984, p. 21) calls it the reflection step that converts experience into skill.

Running this kind of session is not the same as teaching. Delivering content takes little time. What takes time is managing the gap between what a participant attempts and what the LLM produces. Civil servants who are not used to iterating will often conclude after one failed prompt that the tool does not work. The trainer's job is to keep them going. This requires familiarity with LLM prompting and the patience to let each participant reach the result independently. A workshop where the trainer builds the tool and participants watch is a demonstration, not a success experience.

The output of a well-designed session is not a certificate or a slide deck. It is a working artefact that staff will actually use. A participant who leaves knowing why the tool runs locally, why no citizen data went into the LLM, and how to update it when the fee schedule changes, has built the kind of transferable reasoning the AI literacy frameworks in Section 3 call for.

5. CONCLUSION

This paper set out to identify the conditions under which AI literacy training programs lead to sustainable, transferable competence in German municipal administration. The answer, derived from the theoretical frameworks discussed in Section 3 and illustrated by a concrete workshop case, can be stated as a single thesis: sustainable AI competence does not arise from tool instruction, but from situated experiential learning that changes the way participants think about their work.

Section 2 showed that German municipal administrations face a convergence of pressures (labour shortage, legal compliance requirements, and fragmented IT infrastructure) that prevent straightforward AI adoption. Training does not make these pressures go away. But a program that accounts for them will produce more realistic results than one that ignores them.

Section 3 established that lasting adult competence requires more than information transfer. Knowles (1980) showed that adults learn through problems, not content. Mezirow (1997) showed that lasting learning requires a change of perspective, not just the acquisition of new skills. Lave and Wenger (1991) argued that competence is inseparable from the context in which it develops. Kolb (1984) provided a structural model for how experience consolidates into transferable skill. Bandura (1977) explained why a concrete success experience is the most reliable mechanism for building the self-efficacy that sustains new behaviour after training ends. Long and Magerko (2020) and Laupichler et al. (2022, p. 1) confirmed that AI literacy in professional contexts remains underspecified and requires frameworks that prioritise transferable competence over tool-specific knowledge.

From these perspectives, Section 4 derived four conditions: training must be grounded in a real work context; participants must experience AI as a tool to build tools rather than a service to consume; every participant must leave with a working result; and data protection must be an integrated design principle, not a constraint imposed after the fact. The Empfingen workshop illustrated how said conditions can work together in practice: civil servants with no technical background produced a working, compliant, locally-running tool in a single afternoon.

Three limitations should be named. First, the Empfingen case is a single illustration, not an empirical study. The conditions derived here are theoretically grounded but have not been systematically tested across multiple administrations or contexts. Future research should examine whether they hold across different task types, team compositions, and levels of prior AI experience. Second, this paper focuses on the design of individual training interventions. It does not address the organisational conditions that must be in place for trained individuals to actually use what they have learned: management support, peer validation, access to appropriate tools, and time to apply new competences. These factors are likely to determine whether the effects of even well-designed training are sustained at scale. Third, there is the lifecycle of AI-generated tools: a locally built calculator works until the underlying fee schedule changes. At that point, the tool needs to be updated. Staff who built it with LLM assistance once can likely rebuild it, but this dependency is not cost-free and should be factored into any deployment decision.

The broader implication for German municipal administration is this: Knowing about AI and being able to use it are two different things. Closing that gap is a design problem and not an information problem. The tools exist, and the legal framework can be navigated under appropriate design conditions. The remaining gap is therefore not only technical or regulatory, it is also pedagogical.

6. Declaration on the Use of AI-Supported Tools

In accordance with the IU Guideline for the Use of AI-Supported Systems in Studies (Version 1.0, 28 June 2025), the author declares that AI-supported writing assistants were used during the preparation of this paper to assist with formulation and language revision. All content, arguments, source selections, and final editorial decisions reflect the author's own intellectual work.

7. REFERENCES

Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84(2), 191–215. <https://doi.org/10.1037/0033-295X.84.2.191>

Carolus, A., Koch, M. J., Straka, S., Latoschik, M. E., & Wienrich, C. (2023). MAILS – Meta AI literacy scale: Development and testing of an AI literacy questionnaire based on well-founded competency models and psychological change- and meta-competencies. *Computers in Human Behavior: Artificial Humans*, 1, 100014. <https://doi.org/10.1016/j.chbah.2023.100014>

Datenschutzkonferenz (DSK). (2024). Orientierungshilfe zum Einsatz von KI-Systemen durch öffentliche Stellen. DSK. https://www.datenschutzkonferenz-online.de/media/oh/20240506_DSK_Orientierungshilfe_KI_und_Datenschutz.pdf

dbb Beamtenbund und Tarifunion. (2024). Monitor öffentlicher Dienst 2024. dbb Beamtenbund und Tarifunion. https://www.dbb.de/fileadmin/user_upload/globale_elemente/pdfs/2024/dbb_monitor_oeffentlicher_dienst_2024.pdf

European Commission. (2025). Digital Decade 2025 country reports: Germany (SWD(2025) 294 final, Part 11/27). European Commission. <https://digital-strategy.ec.europa.eu/en/factpages/germany-2025-digital-decade-country-report>

European Parliament and Council of the European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). *Official Journal of the European Union*. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act). *Official Journal of the European Union*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689

Kawik, M., Adena, S., Heimann, H. M., Krause, T., Metzger, P.-S., Pflüger, S., Schäffer-Külz, U., & Werner, A. (2022). *Personalgewinnung, -bindung und -entwicklung im öffentlichen Dienst* (1. Aufl.). Luchterhand Verlag.

Knowles, M. S. (1980). *The modern practice of adult education: From pedagogy to andragogy* (2nd ed.). Cambridge Book Company.

Knowles, M. S., Holton, E. F., III, & Swanson, R. A. (2005). *The adult learner: The definitive classic in adult education and human resource development* (6th ed.). Elsevier.

Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development*. Prentice-Hall.

Laupichler, M. C., Aster, A., Schirch, J., & Raupach, T. (2022). Artificial intelligence literacy in higher and adult education: A scoping literature review. *Computers and Education: Artificial Intelligence*, 3, 100101. <https://doi.org/10.1016/j.caeai.2022.100101>

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). ACM. <https://doi.org/10.1145/3313831.3376727>

Mergel, I. (2021). Digital transformation of the German state. In S. Kuhlmann, P. Lorenz-Meyer, & F. Sager (Eds.), *Public administration in Germany* (pp. 331–347). Palgrave Macmillan. https://doi.org/10.1007/978-3-030-53697-8_19

Mezirow, J. (1978). Perspective transformation. *Adult Education Quarterly*, 28(2), 100–110. <https://doi.org/10.1177/074171367802800202>

Mezirow, J. (1995). Transformation theory of adult learning. In M. R. Welton (Ed.), *In defense of the lifeworld* (pp. 39–70). SUNY Press.

Mezirow, J. (1997). Transformative learning: Theory to practice. *New Directions for Adult and Continuing Education*, 74, 5–12. <https://doi.org/10.1002/ace.7401>

Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>