

Are Algorithm Registers Transparent? Perspectives from Germany

Iman Peljto^{1,2}, Xenia Heilmann², Mattia Cerrato²

¹TU Darmstadt
Germany

² Johannes Gutenberg University Mainz
Germany

iman.peljto@stud.tu-darmstadt.de, xenia.heilmann@uni-mainz.de, mcerrato@uni-mainz.de

Abstract

Algorithm registers are public-facing databases that display basic information about algorithms employed in public administration. While several such registers exist across Europe and globally, their capacity to deliver meaningful transparency remains contested. In Germany, the landscape is notably fragmented: no federal-level register exists, yet at least five state- and federal-level initiatives publish information about AI systems with varying scopes and objectives. A recent conceptual proposal by Lorenz (2025) outlines technical and governance requirements for a national AI transparency register in Germany. We repurpose this proposal as an audit instrument, extracting structured checklists from the transparency goals and subgoals formulated by Lorenz (2025). The resulting checklists, translated from German into English, are made publicly available to support practitioners auditing existing registers or designing new ones. We apply this framework to conduct an external audit of the two main existing German transparency initiatives, *Marktplatz der KI-Möglichkeiten* and *Plattform Lernende Systeme*, evaluating the extent to which they fulfill the proposed goals. Our audit reveals that several adaptations are likely needed for these registers to serve as a useful transparency instrument. We further propose a visualization of register transparency levels and derive concrete action items for improving existing German platforms.

Extended version — <https://arxiv.org/abs/2606.02347>

1 Introduction

Algorithm registers are publicly accessible databases documenting algorithmic systems. Murad (2021, p. 16) defines a register as “a log of algorithmic decision making systems used by a public authority that have some level of direct impact on its citizens.” Of particular interest are registers maintained by government bodies: Available registers and standards include the Dutch AI register (Cath and Jansen 2022), the UK (Open Government Partnership 2025), and several others in the EU (Soizic Pénicaud and IA Ciudadana 2025) and globally (Global Partnership on Artificial Intelligence 2024). These national transparency initiatives are increasingly relevant in light of the supra-national EU database for high- risk AI systems listed in Annex III of the EU AI

Act and mandated under Article 71 (European Parliament and Council 2024). Interoperability between national transparency registers and the future EU database could substantially facilitate compliance and reporting obligations across EU member states. In the European setting, the German situation is notable for its fragmentation. While the Bundestag’s Digital Committee has expressed interest in a federal-level register and has conducted a public hearing that touched upon its foundation (David Roth-Isigkeit 2024), we are not aware of a concrete plan towards the establishment of a new register. Rather, there exist at least five state- and federal-level initiatives that deal with the publication of information about AI systems, each with different scopes and objectives (see Section 3.1).

This fragmented landscape has recently received a notable conceptual contribution from Lorenz (2025), who presents a technical and organizational concept for a national AI transparency register in Germany. The project was funded by the VolkswagenStiftung and the work is published by AlgorithmWatch. Lorenz’s report starts from a description of the current limitations in federal-level efforts (Lorenz 2025, pp. 4–6) and proposes an alternative at both the technical and governance levels. In this paper, we repurpose Lorenz’s proposal as an audit instrument: rather than using it to build a new register, we extract from it a checklist of goals and subgoals and apply it to two existing German platforms, *Marktplatz der KI-Möglichkeiten* (MaKI) and *Plattform Lernende Systeme* (PLS) (described in Section 4). Our purpose is to evaluate the *transparency* of existing *transparency initiatives* in Germany. Our main finding is that several gaps exist across multiple dimensions, particularly at the level of training data documentation and risk assessments. To summarize, our main contributions are as follows:

1. We review the European algorithm register literature and organize it along three dimensions of concern (scope, information depth, and effectiveness) that inform our subsequent audit framework.
2. We develop an audit approach structured around the transparency goals suggested by Lorenz (2025) and translate the resulting checklists from German into English (see Appendix B¹), facilitating reuse by practitioners aiming to audit existing registers or develop new

¹See the extended version: <https://arxiv.org/abs/2606.02347>.

ones.

3. We apply this approach to conduct an external audit of MaKI and PLS, analyzing the extent to which the primary and secondary goals of Lorenz (2025, Table 39) are fulfilled.
4. We propose a visualization of register transparency levels and identify concrete action items towards greater fulfillment of transparency goals.

2 Algorithm Registers in Europe: Scope, Standards, and Limitations

Among the transparency instruments available to governments, algorithm registers occupy a distinctive position. Soizic Pénicaud and IA Ciudadana (2025) classify them as *systemic, supply-driven* instruments: proactive disclosures that aim to make an entire portfolio of systems visible, shifting the burden of initiative from the public to the institution.

Amsterdam and Helsinki became the first cities to publish such registers in September 2020 (Floridi 2020; Haataja, van de Fliert, and Rautio 2020). Adoption has since accelerated: the Eurocities network developed a shared Algorithmic Transparency Standard involving nine cities (Eurocities Network 2026); the UK launched its Algorithmic Transparency Recording Standard (ATRS) in 2021 and made it mandatory for all central government departments in 2024 (Open Government Partnership 2025; Leslie and Selman 2024); and the Netherlands launched a national register (algoritmes.overheid.nl) in December 2022. By 2024, the Global Partnership on AI had mapped 69 active repositories worldwide (Global Partnership on Artificial Intelligence 2024), while Soizic Pénicaud and IA Ciudadana (2025) counted 34 registers in the EU alone. At the EU level, the AI Act (European Parliament and Council 2024) mandates a public database of high-risk AI systems (Article 71), though systems used in law enforcement, migration, asylum, and border control are registered in a non-public section (Article 49(4)).

Despite this rapid proliferation, the literature reveals recurring concerns about the design and effectiveness of these instruments. We organize these concerns along three dimensions: the *scope problem* (what gets registered), the *information problem* (what is disclosed about each registered system), and the *effectiveness problem* (whether registers achieve their stated goals in practice).

2.1 The Scope Problem: What Gets Registered?

A foundational design question for any register is what counts as an “algorithm” or “AI system” that should be registered. The AI Act’s definition (Article 3(1)) centers on machine-based systems that operate with varying levels of autonomy and that infer how to generate outputs from their inputs. The European Commission published non-binding guidelines in February 2025 attempting to clarify the boundary, citing linear and logistic regression as explicit examples of non-AI systems (European Commission 2025, Paragraph 46) while including symbolic methods such as expert systems (European Commission 2025, Paragraph 39). The

guidelines ultimately acknowledge that “no automatic determination or exhaustive lists of systems that either fall within or outside the definition of an AI system are possible” (*ibid.*, Paragraph 62). Soizic Pénicaud and IA Ciudadana (2025) flag this ambiguity as a direct concern for registers: if the definition excludes rule-based automation, then systems that may carry substantial risks, such as deterministic scoring or eligibility algorithms, would fall outside the scope of both the EU database and any national register that adopts the AI Act definition. Murad (2021, pp. 6–7, 18–19) frames the issue more broadly as a question about the appropriate “unit of disclosure,” arguing that transparency should apply to the algorithmic decision-making system as a whole rather than to the algorithm in a narrow technical sense.

Beyond the definition of what *qualifies* for registration, there is the question of what *actually gets registered* in practice. Cath and Jansen (2022, Sec. 1–2) show how the Amsterdam and Helsinki registers were narrowly scoped to cover “a set of largely uncontentious bureaucratic municipal uses of AI” while excluding sensitive domains such as welfare fraud detection and predictive policing. They argue that these selective inclusions and exclusions are themselves political acts. Nieuwenhuizen (2024, p. 418) generalizes this observation, noting that the unit of disclosure is “inherently a political choice, impacting the register’s comprehensiveness and reflecting the organization’s values and priorities.” Her empirical study found that more than half of Dutch municipalities had registered nothing at all. A further scope question is whether to include systems still in development or already decommissioned: the Dutch standard permits but does not require pre-deployment registration; the Eurocities standard provides a seven-stage lifecycle from *planned* through *retired*; and Soizic Pénicaud and IA Ciudadana (2025) and the Ada Lovelace Institute et al. (Ada Lovelace Institute, AI Now Institute, and Open Government Partnership 2021) advocate for pre-deployment registration as a condition for early citizen participation.

2.2 The Information Problem: What Is Disclosed?

Even when a system is registered, the question remains: what information is captured, and in what form? The two most developed European standards illustrate the range of approaches. Other national registers, such as Scotland’s AI Register² and Norway’s public-sector AI overview,³ are not analyzed separately here: Scotland’s register is built on the same technical platform used by the Amsterdam and Helsinki registers (Haataja, van de Fliert, and Rautio 2020) and does not define an independent schema; Norway’s overview is a catalog of projects within a broader national data portal, without published field definitions that could be compared against the Dutch or Eurocities schemas.

The **Dutch National Publication Standard** (v1.0.10) (Ministerie van Binnenlandse Zaken en Koninkrijksrelaties 2024) organizes information into four sections: General Information, Responsible Use, Functioning, and Metadata. It

²<https://scottishairegister.com>

³<https://data.norge.no/kunstig-intelligens>

contains no dedicated training data field; the closest approximation is the free-text “Gegevens” (Data) field, which does not distinguish training data from operational input and imposes no structure on data provenance, collection period, or bias assessment. The “Technische werking” (Technical functioning) field must indicate whether the system is self-learning but provides no further sub-fields. The stated objective is to protect “people, their fundamental rights, and public values” and to “enable citizens to critically monitor and question the government.”⁴

The **Eurocities Algorithmic Transparency Standard** (v0.3.1) (Eurocities Network 2026) is more ambitious. It introduces `training_data` as a named, required field separate from `source_data` (operational input). An optional `data_bias` field explicitly asks whether data contain biases and how these are compensated for. The Oversight category mandates documentation of performance monitoring, human intervention procedures, type of use (descriptive, diagnostic, predictive, or prescriptive), and objection mechanisms. However, all fields remain free-text prose without structured sub-fields, and the “required” designation carries no enforcement since participation is voluntary. Leslie and Selman (2024), comparing transparency requirements across five jurisdictions, find that no country has yet established legally enforceable, comprehensive disclosure standards for public-sector AI, reinforcing the concern that even the most developed register standards remain voluntary instruments.

A notable gap across both standards, and across the German proposal discussed in Section 3, is the absence of any field for attaching or referencing model cards (Mitchell et al. 2019) or datasheets for datasets (Gebru et al. 2021). Including such a mechanism would enable layered disclosure, where a general audience reads the register entry while a more technical audience can access structured, model-level documentation.

2.3 The Effectiveness Problem: Do Registers Work?

The most fundamental criticism of algorithm registers is that their existence may create a misleading impression of accountability. Cath and Jansen (2022, Sec. 4) warn against “romanticizing the light touch ex-post transparency measures” provided by registers, arguing that academic enthusiasm about their mere existence can produce what they term “ethics theater”: visible but substantively hollow measures that create a false sense of oversight. Busuioc, Curtin, and Almada (2023) raise a related concern in their analysis of the EU AI Act: substituting transparency with mediated explanations, where the provider controls what is made visible, opens disclosure to manipulation. In a register context, this applies directly, since the registering body is typically also the deployer and no independent verification is required.

A complementary line of criticism asks: *transparent to whom?* Kemper and Kolkman (2019) argue that transparency only becomes accountability when a *critical audience* exists that is capable of interpreting and acting on the

disclosed information. Bell, Nov, and Stoyanovich (2023) build on this reasoning, proposing that disclosure mechanisms should be designed based on what different audiences need. Nieuwenhuizen (2024) provides the most direct empirical evidence: oversight authorities and civil society watchdogs in the Netherlands reported that the information disclosed was of limited usefulness, and only 5% of entries included a fundamental rights assessment. She concludes that transparency intermediaries should be the targeted audience if external accountability is the goal.

Yet Nieuwenhuizen’s findings are not uniformly negative. She also documents a *disciplining effect*: the process of registration compels civil servants to reflect on whether a system is justified, what its legal basis is, and how risks are managed. Her characterization of registers as a “meaningful box-ticking exercise” captures the duality: the exercise creates organizational value even when its public-facing output is thin. Similar findings from six case studies at Dutch public organizations are reported by van Vliet et al. (2024), who propose a reference process for register implementation. Soizic Pénicaud and IA Ciudadana (2025) distill from their survey of 34 EU registers a set of design recommendations, including centralized mandatory registration, collaboration with civil society in defining categories, coverage of in-development and decommissioned systems, and a requirement that registering bodies document and justify the absence of information rather than leaving fields empty.

The recurring finding across this literature is that registers function better as internal governance tools, prompting reflection and organizational learning, than as genuine public accountability mechanisms enabling external scrutiny. This gap between aspiration and practice motivates the present paper: rather than proposing yet another standard, we take a recent comprehensive proposal for a German national register (Lorenz 2025) and derive from it a checklist of goals, which we then use to evaluate the existing German initiatives.

3 Proposal for a German National AI Register

Lorenz (2025) presents a detailed concept for the institutional, organizational, and technical implementation of a German national AI transparency register (KI-Transparenzregister). It is, to our knowledge, the most comprehensive proposal for a German register, covering not only the information schema but also governance structures, technical architecture, and stakeholder engagement. For this reason, we use it as the basis for our evaluation framework. We summarize its main features below before describing, in Section 5, how we transform its goals and subgoals into a checklist for evaluating existing registers.

The report is written in German. The German institutional names and technical terms used below follow Lorenz’s original text; English translations are ours unless otherwise noted.

3.1 Scope and Design Choices

Lorenz (2025) begins by surveying the fragmented German

⁴<https://www.nldigitalgovernment.nl/overview/algorithms/>

landscape of AI-related platforms, identifying five federal- and state-level initiatives (described in Section 4): MaKI, PLS, the KI-Observatorium, Mission KI, and KI.NRW. These differ substantially in scope, audience, and standards, and none functions as a comprehensive public transparency register. Furthermore, since December 2025, Mission KI is no longer operating. While Lorenz does not offer a full evaluation of MaKI, she comments on the lack of obligatory reporting and the lack of coordination between platforms (Lorenz 2025, pp. 4–5): the same AI applications sometimes appear on multiple platforms under different names or with divergent descriptions.

From this assessment, Lorenz defines five target groups (citizens, public authorities, supervisory bodies, civil society and academia, and private-sector organizations) and five primary goals: safeguarding fundamental rights, strengthening public trust, enabling dialogue and review, centralizing AI system information, and ensuring interoperability with the EU high-risk AI database. Three secondary goals concern knowledge sharing, cross-sector cooperation, and minimizing administrative burden. The interoperability goal is notable: Lorenz’s proposal is designed as a national complement to the EU database mandated by Article 71 of the AI Act, intended to go beyond it in both coverage (including systems below the high-risk threshold) and information richness.

Several further design choices are worth highlighting. First, Lorenz aligns the definition of AI systems with the EU AI Act, arguing that this facilitates compliance, comparability, and legal certainty. Second, she restricts coverage to systems that are or have been deployed (“produktiv”), explicitly excluding systems in planning or development. Her rationale is to maximize public relevance and avoid the problem she observes in MaKI, where only a quarter of listed systems were operational⁵. This stands in contrast with the Dutch register (which permits pre-deployment registration), the Eurocities standard (which provides a full seven-stage lifecycle), and the recommendations of both Soizic Pénicaud and IA Ciudadana (2025) and the Ada Lovelace Institute et al. (Ada Lovelace Institute, AI Now Institute, and Open Government Partnership 2021), who advocate for pre-deployment registration as a condition for early citizen participation. Third, Lorenz explicitly declares that the register should not itself serve as a mechanism for reviewing ethical, legal, or data protection compliance; it should document such reviews, not conduct them. Her own schema, however, includes nine fields dedicated to risk classification, risk assessment results, and impact assessment outcomes, creating a tension between documenting compliance outputs and disclaiming responsibility for their quality.

3.2 Information Schema

The proposed schema contains 30 fields, divided into externally visible and internal fields, organized across four categories: information about the responsible institution (Informationen zur verantwortlichen Institution), information

about the AI system (Informationen zum KI-System), information on criticality (Informationen zur Kritikalität des KI-Systems), and metadata about the submitting person (Angaben zur übermittelnden Person). We provide the translated field names and descriptions in the Appendix, Table 3.

The treatment of training data is minimal. Field 16 calls for a “basic and concise description of the information used and the operating logic” (transl. from “Grundlegende und knappe Beschreibung der verwendeten Informationen (Daten, Eingänge) und der Betriebslogik”), citing Articles 51 and 53 of the AI Act. This field does not distinguish between training data and operational input, and provides no structured sub-fields for dataset provenance, size, collection period, or bias assessment. In this respect, the proposal is comparable to the Dutch national standard and less detailed than the Eurocities standard, which mandates separate `training_data` and `source_data` fields and offers an optional `data_bias` field. Field 16 does include a provision for general-purpose AI systems, where the free-text description may be replaced by an upload of a training data summary template (“Zusammenfassung der Trainingsdaten”).

The criticality section is more developed, comprising nine fields covering risk classification under the AI Act, risk assessment, identified risks, mitigation measures, data protection impact assessment, and fundamental rights impact assessment. Across these fields, the same logic is applied: one field documents whether and how an assessment was conducted, while the corresponding field captures its results. Lorenz recommends following established risk assessment frameworks (ISO 31000:2018, ISO/IEC 23894:2023, NIST AI RMF, or IEEE EAD) and connecting register entries to the risk identification and evaluation steps prescribed therein.

3.3 Architecture, Governance, and Outreach

A distinguishing feature of the proposal is its depth on technical and institutional matters. On the technical side, Lorenz describes distinct user interfaces for data entry, analysis, reporting, and feedback, and devotes attention to what she terms *bitemporale Historisierung* (bitemporal historization), the time-stamping and versioning of all register data. The reporting function envisions individual AI system profiles (KI-Steckbriefe), geographic maps of AI deployment (KI-Landkarten), aggregated statistics, and an API supporting automated analyses and integration with other databases. Any user of the register may submit feedback, which is connected to an evaluation of the register’s own impact.

On the governance side, Lorenz proposes a tripartite division of responsibility: professional (setting guidelines and monitoring requirements), organizational (coordinating activities and stakeholder relations), and technical (infrastructure operation, potentially outsourced). She notes ongoing discussions about designating the Bundesnetzagentur⁶ as the

⁵As of May 2026, this still holds. Currently, only two-thirds of MaKI’s entries are operational.

⁶The Bundesnetzagentur für Elektrizität, Gas, Telekommunikation, Post und Eisenbahnen is Germany’s multi-sector infrastructure regulator. It does not offer a translation of its name to English (see <https://www.bundesnetzagentur.de/EN/>).

coordinating body for national AI supervision, with a new dedicated digital agency as an alternative.

Finally, Lorenz integrates information campaigns, workshops, best practice templates, and citizen forums (Bürgerforen) into the register concept itself, treating outreach and capacity building as a component of the register rather than an afterthought. This echoes Soizic Pénicaud and IA Ciudadana (2025)’s recommendation for active dissemination and responds to the empirical finding that registers require informed audiences to produce accountability effects (Kemper and Kolkman 2019; Nieuwenhuizen 2024).

3.4 From Proposal to Evaluation Framework

Lorenz’s report is a normative proposal rather than an empirical study. Its value for our purposes lies not in whether its recommendations are optimal, but in the fact that it articulates, in concrete and operational terms, what a comprehensive German AI register should contain and how it should function. By taking the proposal at face value and extracting from it a set of (sub)goals, we obtain a checklist that can be applied to existing German initiatives to assess the extent to which they fulfill their stated transparency objectives. This approach has limitations: the checklist inherits any blind spots in Lorenz’s proposal, and it evaluates registers against their purported goals rather than against an ideal standard. We accept these limitations deliberately. We aim to provide a structured, reproducible method for evaluating the registers that exist.

4 Existing AI Platforms in Germany

Having summarized Lorenz’s proposal and her assessment of the current landscape, we now describe the two platforms that form the subject of our audit.

4.1 Marktplatz der KI-Möglichkeiten

The “Marktplatz der KI-Möglichkeiten” (MaKI)⁷ describes itself as a central AI transparency register for the federal public administration. According to its official mission, the platform aims to provide a comprehensive overview of AI use across the German public administration, facilitate exchange on AI systems and applications, and promote the reuse of existing solutions (Das vom BMI initiierte Beratungszentrum für Künstliche Intelligenz (BeKI) 2025). It provides detailed and accessible information on individual AI applications, referred to as *KI-Steckbriefe* (AI fact sheets). Launched in January 2025, MaKI was conceived as a central coordination and networking tool for AI use within the federal administration. At its introduction, the platform listed only forty-six productive AI applications, with participation remaining voluntary. As of May 2026, 115 active systems are described, with 205 more at various stages of planning and decommissioning. The model of registration is voluntary: as such, our evaluation is that rather than promoting transparency to the public, MaKI primarily serves internal purposes, connecting administrative bodies dealing with AI systems and supporting internal knowledge exchange.

⁷<https://maki.beki.bund.de>

4.2 Plattform Lernende Systeme (Platform Learning Systems)

The “Plattform Lernende Systeme” (PLS)⁸ is an initiative funded by the Federal Ministry of Education and Research (BMBF) of Germany. It serves primarily as an exchange and dialogue platform for AI research, innovation, and application development in Germany. Its focus is on the scientific and industrial advancement of AI, not on governmental or administrative use, which distinguishes it clearly from initiatives like MaKI. PLS collects and presents examples of AI systems developed or deployed in Germany, emphasizing innovation transfer, best practices, and cross-sector learning. Beyond individual case studies, the platform also provides a searchable and filterable overview of registered AI systems, which can be explored either through an interactive *KI-Landkarte* (AI map) or in a structured list view. The *KI-Landkarte* visualizes the geographical distribution of AI systems across Germany, helping to identify regional concentrations, innovation clusters, and areas of potential policy or research interest. At the time of data extraction, the platform contained 1,305 registered AI projects in Germany, together with structured metadata on their application domain, technology field, industry sector, funding context, and developer information.

4.3 Other Platforms

Beyond MaKI and PLS, several smaller platforms serve various purposes, namely the KI-Observatorium, Mission-KI, and KI.NRW. According to Lorenz (2025), KI-Observatorium and Mission-KI cannot be regarded as initiatives on the road to such a register. Nevertheless, both platforms provide useful products, information, and statistics on AI-related topics for the general public. For example, Mission-KI has developed a quality standard for low-risk AI, but has stopped operation in December 2025 after a change in the German government. KI-Observatorium has published articles on AI trends and terminology. KI.NRW offers an AI map of applications in North Rhine-Westphalia; however, because it covers only a single federal state, it does not provide information at the national level. Moreover, as the platform is essentially a map, populating the tables required for this analysis yields few entries. Consequently, MaKI and PLS remain the two primary platforms examined in this paper.

5 Methodology

Lorenz’s concept paper was written as a development guide: it specifies what a future German register should contain and how it should be built. We repurpose it as an audit instrument, using its requirements not to build a register but to evaluate existing ones. This shift in perspective is deliberate. The development guide articulates what *should* be present; an audit checks what *is* present. The gap between the two reveals the distance that existing platforms must travel to meet the transparency goals of algorithm registers. This reorientation introduces a structural limitation: some of Lorenz’s

⁸<https://www.plattform-lernende-systeme.de>




Figure 1: Methodology overview.

subgoals presuppose access to internal architecture, governance processes, or implementation decisions that are not visible from outside. Rather than discarding these items, we retain them and classify them as *No Access*, so that the limits of what an external audit can assess become visible in the results. This visibility can itself support accountability, particularly in cases where restricted access may not be clearly justified by privacy, security, or safety concerns. We proceed in five steps (Figure 1).

5.1 Step 1: Review

As a first step, we review and translate the concept paper by Lorenz (2025) into English. During this process, we analyze each section individually to identify its core objectives and the normative framework it establishes for transparency registers. We also summarize the overarching requirements articulated throughout the paper. The resulting synthesis forms the basis of Section 3 and provides the conceptual foundation for all subsequent steps.

5.2 Step 2: Audit Checklists

Although the concept paper by Lorenz (2025) was originally conceived as development guidelines for a national transparency register, we repurpose its content for the context of auditing existing registers. This shift in perspective requires translation of prescriptive design criteria into audit items which can be evaluated empirically. For each section of the concept paper, we identify the overarching requirements and supporting subgoals needed to achieve the intended goals, then translate them into structured checklist items. This preserves the original structure of the paper while enabling systematic comparison across platforms. Where requirements are formulated abstractly, we operationalize them into concrete indicators to the extent that the source material permits. The resulting checklists are provided in English in the Appendix B, Table 3 to 9 with each table corresponding to one chapter of the underlying concept paper. In addition, we extract the primary and secondary goals of a transparency register as a whole, as defined in Lorenz (2025, Table 39), and render these in checklist format as well.

5.3 Step 3: Platform Assessment

We apply the checklists developed in Step 2 to the two German platforms under review: MaKI and PLS. For each checklist item, we compare the requirements defined by

Lorenz (2025) against the information publicly available on each platform, drawing exclusively on content accessible to any external evaluator without special authorization or technical access. It should be noted that a recurring obstacle throughout this assessment is restricted access to certain data. This is a structural consequence of our methodological reorientation: we employ the proposal by Lorenz (2025) as an auditing instrument, whereas it was originally designed as a development guide. As a result, some requirements presume insider knowledge of system architecture, internal processes, or implementation decisions that MaKI and PLS do not disclose publicly. Each checklist item is classified into one of four categories: *Yes*, *No*, *Partially*, and *No Access*. Together, these classifications yield a comprehensive picture of which requirements are satisfied, partially satisfied, or impossible to assess for the public. The full results are documented in the Appendix B Table 3 to 9.

Yes A criterion is marked *Yes* when the required information or feature is directly locatable and observable on the platform in a form that fully satisfies the stated requirement. This classification is applied only to criteria that can be evaluated on the basis of observable content alone, i.e., the presence of a system name, provider information, a description of the AI system's purpose, or publicly accessible explanatory material.

No A criterion is marked *No* when the required information or feature is absent from the platform, and where its presence would reasonably be expected to be publicly verifiable. This classification signals a documentable gap as the platform does not indicate that the requirement is addressed anywhere on its publicly accessible interface.

Partially *Partially* fulfilled denotes the classification applied when a criterion is observable on the platform and shows evidence of implementation, but fails to meet the full scope of the goal due to incomplete execution, missing components, or functionality that cannot be verified in its entirety. This category captures insufficient compliance: the platform has engaged with the requirement but has not fully discharged it. For instance, MaKI includes dedicated fields for data protection impact assessments under GDPR Article 35, yet these fields remain unpopulated for all analyzed AI systems, and no public justification is provided for the absence of this information. The infrastructure for compliance exists; the information does not. The prevalence of only

partial implementation therefore raises questions about how meaningful this transparency is in practice (see Section 2). Crucially, this classification is restricted to cases where empirical evidence of partial implementation is directly observable. It is distinguished from *No* where no relevant feature or attempt at implementation exists. For all entries classified as *Partially*, we provide an explanation in Appendix A.1.

No Access *No Access* denotes the classification applied when a criterion cannot be empirically evaluated because the necessary information is not publicly observable or technically verifiable by an external reviewer. Importantly, this classification does not imply non-compliance, but rather reflects that the required information remains inaccessible to external auditors. In practice, this category arises in two main contexts. First, it applies to internal system architecture such as encryption protocols, authentication mechanisms, or identity and access management procedures. For these, neither platform provides publicly accessible technical documentation. Second, it applies to criteria that presuppose external institutional infrastructure, such as interoperability with an up-to-now unavailable EU-wide AI database, which falls outside the scope of any individual platform's current capabilities. The entries labeled *No Access* are described in detail in Appendix A.2.

5.4 Step 4: Evaluation

In the fourth step, we relate the item-level findings from Step 3 to the main and secondary (sub)goals for German transparency registers outlined by Lorenz (2025, Table 39). For each subgoal, we aggregate the classifications of all corresponding checklist items and visualize the distribution across the four categories using a color-coded pie-chart: **Yes**, **No**, **Partially**, and **No Access**. Concretely, each checklist item that maps onto a given goal contributes one observation to the distribution. The concrete mapping of which checklists and items contribute to which goals can be found in Appendix B Table 2. The resulting pie chart thus reflects not only whether a goal is met, but to what extent and with what degree of consistency across the individual criteria that operationalize it. A goal for which most items are classified *Yes* can be considered substantially fulfilled, whereas a goal dominated by *No* or *No Access* points to a structural gap in the platform's implementation. This aggregation allows us to move beyond item-level compliance and assess the degree to which each platform fulfills the broader functions that a transparency register is expected to serve. The aggregated results for the main goals are presented in Table 1 and discussed in detail in Section 6. A presentation of the secondary goals can be found in Table 10 in the Appendix.

5.5 Step 5: Action Items

Building on the evaluation in Step 4, we derive a set of platform-specific action items intended to move items currently classified as *No* or *Partially* toward full compliance. Rather than issuing generic recommendations, each action item is grounded directly in the gap analysis produced by Steps 3 and 4 and is tailored to the specific deficiencies

identified for MaKI and PLS, respectively. Where Step 4 establishes *that* a gap exists, Step 5 addresses *how* it might be closed. The resulting recommendations serve a dual purpose. For platform operators, they constitute an actionable roadmap: each item identifies an addressable deficiency and specifies the change required to resolve it. For policymakers and civil society actors, the action items collectively function as a diagnostic instrument, surfacing the structural patterns, such as a recurring absence of technical metadata or a systematic gap in risk assessment, that individual item-level findings alone might obscure. It should be noted that the action items are bounded by the information available through the methodology described in Steps 1-4. We do not focus on deficiencies which stem from upstream factors, such as the absence of a mandatory reporting obligation or the lack of standardized taxonomies across the German public sector. The action items are presented alongside the evaluation results in Table 1 and elaborated upon in the next section. A complete list is also available in the Appendix D.

6 Findings

We present our compliance visualization of the primary goals in Table 1. The visualization is structured so that pie charts are available for all primary goals and subgoals of a register as conceptualized in this paper and by Lorenz (2025). The rationale for this presentation is to provide a compact overview of how many checklist items are fulfilled, partially fulfilled, not fulfilled or currently not assessable due to our access limitations as external auditors.

We note that Table 1 does not provide a comparative weighting of the checklist items. Thus, the resulting visualization should not be taken as a complete assessment of the level of fulfillment of certain goals. To give a concrete example, two of the unfulfilled checklist items in subgoal 1a (Publication of detailed information on purpose, function, benefits, risks, legal basis, and responsible institution) for MaKI are i) to offer the responsible institution's logo on the AI system's description page; ii) discussing the fundamental rights impact assessment results as per Article 27(1)(a)–(f) of the EU AI Act. Clearly, these two points do not warrant the same urgency, but take equal space in the visualization in Table 1. The purpose of the present section and analysis will therefore be to further qualify the data we extracted so to present the main results of our audit.

Goal 1: Ensuring transparency regarding deployed AI systems In terms of the published basic information about the system (**Goal 1a**), MaKI and PLS perform reasonably well; the picture differs when one moves onto information about criticality and risks of the system. PLS does not publish any information about the risk classification of the system, the risk assessment or identified risks. MaKI does have this information, but the treatment of risk assessment is minimal, offering a classification but no further details. Compared to Lorenz's proposal, a URL to a document providing concrete, detailed information about risk assessment methodology employed and the result of the analysis is not included. In terms of whether the two platforms review the submitted data (**Goal 1b**), we can provide only a limited as-

Goals	MaKI		PLS	
1: Ensuring transparency regarding deployed AI systems				
1a: Publication of detailed information on purpose, function, benefits, risks, legal basis, and responsible institution	MA1-3, MA7, B1-3			LS1-7, B1-3
1b: Automated and manual review of the submitted data for consistency, completeness, clarity, and timeliness	B13, B14			B13, B14, LS9
1c: Provision of easily accessible, understandable data for the public and capacity-building for its use	B4-12			B4-12, LS9
2: Strengthening trust in the work of public administration				
2a: Identification of fields of application, criticality, and responsible actors	MA1-3, MA7, B1-3, B16			LS1-7, LS9, B1-3, B16
3: Enabling dialogue and oversight				
3a: Publication of reports to support citizen participation and implementation of feedback mechanisms for citizens and organizations	MA5, MA6, B11, B12, B14, B15, B17, B18			B11, B12, B14, B15, B17, B18
3b: Implementation of traceability and reporting obligations	MA5, MA6, B11, B12, B14, B16-18			LS8, LS9, B11, B12, B14, B16-18
3c: Awareness-raising through network building, citizen forums, and long-term partnerships	B4-8			B4-8
4: Centralization of information on AI systems				
4a: Establishment of a central register documenting all AI systems	B13			B13, LS9
4b: Development of an interoperable architecture with automated interfaces	B13			B13, LS9
5: Ensuring interoperability with the EU database				
5a: Implementation of an automated interface for communication with the EU database	MA1-3, MA7, B1-3, B13, B16			LS1-7, LS9, B1-3, B13, B16
5b: Ensuring regular data updates and synchronization	B13, B16, B19			LS9, B13, B16, B19
5c: Substantive consistency of information fields with the EU database and cooperation with EU supervisory authorities	MA1-3, MA7, B1-3, B13			LS1-7, LS9, B1-3, B13

Table 1: Primary goals and subgoals of an AI register, how MaKI and PLS fulfill them, and the action items needed to meet them.

assessment for MaKI, as detailed information is not available to external users, nor were we able to assess the data entry forms, as these are not public. For PLS, an input form is available and we were therefore able to assess some capabilities for automated validation (see Appendix A.1). However, the overall procedures for data quality assurance are not visible to us. In particular, we are not able to assess whether the data is updated every few months and whether the systems appear soon after their commissioning, or whether an update to the data is required after substantial changes. Lastly, both platforms have some ways to support the creation of understandable data for the public (**Goal 1c**). Here, MaKI's exporting capabilities and search functionalities stand out, while PLS offers a map-based visualization of which systems are active in different states as a landing page of sorts. Capacity-building is perhaps the main weakness here, as there are only limited feedback mechanisms for both platforms. Compared to our checklist, there is no citizen forum or an equivalent dialogue channel by which the public may be able to ask questions, raise concerns or provide feedback, nor to signal problems with specific data entries.

Goal 2: Strengthening trust in the work of the public administration We observe here that, at least in name, both MaKI and PLS allow public institutions at the federal, state, and municipal levels to publish information on AI systems, which supports institutional visibility and accountability. MaKI appears somewhat stronger in this regard because it is specifically focused on AI systems used in public administration and therefore aligns more closely with the objective of administrative transparency. However, the platform is not open to external input, which limits broader participation and independent scrutiny. In addition, both initiatives provide only limited information on important trust-related aspects such as risk assessments, data protection impact assessments, and risk mitigation measures. This weakens the ability of citizens and affected groups to evaluate whether AI systems are being deployed responsibly and in compliance with legal safeguards. Finally, stronger alignment with standardized EU reporting structures and the future EU database is needed to improve comparability, consistency, and trustworthiness across public-sector AI registers.

Goal 3: Enabling dialogue and oversight Regarding **Goal 3a**, both MaKI and PLS provide publicly accessible system profiles and geographical visualizations that help citizens identify AI systems active in their municipality or state. However, both initiatives lack sufficiently developed feedback mechanisms that would allow citizens or organizations to submit structured feedback on individual entries or on the platforms as a whole. In addition, the platforms do not explicitly address diverse or critical audiences, and more in-depth information on AI systems, such as training data or validation procedures, is generally not publicly available. Concerning **Goal 3b**, the implementation of reporting obligations and traceability mechanisms remains partial and difficult to evaluate from the perspective of external auditors. PLS appears more transparent regarding submission structures, whereas MaKI restricts data entry to public administration actors. However, both initiatives provide limited pub-

lic information on validation processes, quality assurance, database architecture, and update procedures, and the reporting model currently remains voluntary. With respect to **Goal 3c**, one positive aspect is the establishment of collaborations with educational and civil society organizations and the availability of examples of best practices. However, to the best of our knowledge, no guides or tutorials for using the registers are available, nor have workshops or seminars on reporting been organized. More broadly, there is little evidence of participatory structures such as citizen forums, long-term partnerships, or independent advisory boards integrating ethical, legal, societal, and technical perspectives. Overall, the absence of the AI supervisory authority (*KI-Aufsichtsbehörde*) envisioned by Lorenz (2025) makes institutional responsibilities difficult to assess, but this only increases the importance of transparent reporting and participation mechanisms within the platforms.⁹

Goal 4: Centralization of information on AI systems

Both MaKI and PLS provide web-based interfaces and appear to possess the technical infrastructure necessary to support this objective. At the same time, the assessment of both initiatives remains limited because there is insufficient public information about their database architecture, synchronization interfaces, and quality assurance processes. MaKI is more closely aligned with the requirements of a formal transparency register because it uses standardized fields corresponding to Annex VIII of the AI Act, making it more suitable for structured reporting and legal documentation. In contrast, PLS offers a broader overview of the AI ecosystem, but its data structure is less clearly aligned with the reporting requirements of Annex VIII, particularly regarding conformity assessment documentation such as the EU Declaration of Conformity.

Goal 5: Ensuring interoperability with the EU database

Although Article 71 of the AI Act provides for an EU AI database, it was not publicly operational at the time of our evaluation, making the assessment of interoperability largely prospective. Regarding **Goal 5a**, neither MaKI nor PLS currently provides clearly documented automated synchronization interfaces with the future EU database or with external analysis and reporting tools. However, MaKI appears better positioned in this regard, as mentioned above, because its standardized information fields are more closely aligned with the requirements of the AI Act and Annex VIII. Concerning **Goal 5b**, it also remains unclear whether either initiative supports versioning or bitemporal historization that would make changes to AI system entries traceable over time, including updates, corrections, or changes in deployment and responsibility. With respect to **Goal 5c**, cooperation and synchronization with EU supervisory authorities are not yet established, and the assessment remains limited because the relevant technical details and governance structures are not publicly accessible and supervisory responsibilities are themselves still emerging.

⁹Researchers replicating this analysis for other EU member states should verify whether an equivalent institution exists and whether it has proposed a concrete governance framework.

7 Action Needed and Discussion

Summarizing the preceding section, we find that the goals of a register are only partially fulfilled by the two platforms we evaluated. We give in Appendix D a complete list of the actions needed for these platforms to fulfill the objectives of a register; in the following, we touch upon the most urgent and limiting issues we find in MaKI and PLS.

Transparency towards the public is not the core mission. MaKI's self-described mission as a "marketplace for AI opportunities" undermines its credibility as a neutral transparency instrument. PLS, while broader in scope, similarly lacks a clear mandate to serve as a register upon which accountability claims can be grounded. For either initiative to function as a national transparency register, the primary and overarching goals must be transparency toward the public. Crucially, such a register should be developed with the people it is meant to serve, through participatory processes such as user surveys or public consultations. We find no evidence that either MaKI or PLS has undertaken such an effort, or is planned for the future.

Lack of information on risk and impact assessments. Both registers lack detailed fields for risk and impact assessments. Since neither platform is oriented towards public-facing transparency, serving instead as a promotional instrument for government innovation, any future effort to reach a critical audience (Kemper and Kolkman 2019) will need to collect this information anew, duplicating work that could have been captured at registration. MaKI offers a minimal risk categorization for some systems, but provides no concrete description of oversight measures or actions to be taken should risks materialize. The current registers therefore offer limited transparency (Goal 1) and cannot meaningfully strengthen trust in public administration (Goal 2).

Quality of data entries. Although MaKI uses standardized data formats listed in Annex VIII of the AI Act more consistently than PLS, the content of certain information fields remains questionable. Categories within the MaKI field "AI Method" combine fundamentally different conceptual levels, including methodologies, application domains, and adjacent technical disciplines. Particularly ambiguous classifications such as "Semantic AI," "Statistical Modeling," and "Information Retrieval" illustrate the absence of proper terminology. This undermines the objective of an algorithm register because the purpose of such a register is not merely to label systems as "AI," but to make them interpretable and comparable across institutions and use cases. This may also indicate insufficient quality assurance mechanisms for data entered by reporting authorities and organizations.

Lack of concrete feedback mechanisms. Neither platform offers citizens a way to give feedback on registered systems or pose questions. There is no public dialogue support (e.g. comment fields), no user satisfaction surveys, and no clear consultation opportunities for affected citizens. This impacts Goals 2 and 3, "Strengthening trust" and "Enabling dialogue and oversight."

Public documentation. Neither platform publishes information on its data collection and quality assurance processes, such as entry deadlines, update obligations, or review steps, nor does it provide user guides for navigating and in-

terpreting register entries. MaKI offers export functionality for systematic analysis, but neither platform provides API access.

Coverage. We were unable to confirm whether all relevant AI systems, or only a selection, are documented on MaKI and PLS (Goal 4). Echoing Cath and Jansen (2022), we note the absence of any systems published by the Federal Criminal Police (Bundeskriminalamt) or the Federal Police Office (Bundespolizeiamt), despite credible reports of AI usage in law enforcement¹⁰. Furthermore, in MaKI only AI systems used by the public administration are registered, limiting the coverage by design. We believe that neither initiative operates on a compulsory basis; both rely on voluntary participation, whereas a national register would require well-defined criteria specifying which systems must be registered and which are exempt.

8 Conclusion

This paper examined MaKI and PLS as the two main German initiatives for registering AI systems. Our analysis shows that both contribute to the transparency of AI systems, but only partially fulfill the functions expected from a meaningful national transparency register. The central question is not which information a transparency register claims to capture but whether this information is actually present in the register, how understandable it is, and whether the public can meaningfully use it. Germany currently lacks a centralized national AI transparency register. Existing initiatives remain fragmented, differ substantially in purpose and structure, and appear oriented toward internal administrative documentation rather than meaningful transparency for the public. As a result, important elements such as risk and impact assessments, documentation quality, public explanations, feedback mechanisms, identifiable contact persons for inquiries, and accountability for the accuracy of entries remain underdeveloped. Rather than creating additional disconnected initiatives, future efforts should focus on consolidating and strengthening existing ones, complemented by interviews with operators, administrators, and users to capture practical experiences not visible from the register content alone, and by validating the checklists applied here against other national and international AI registration initiatives. Combining the strongest elements of MaKI and PLS could provide the foundation for a national transparency register that is coherent, publicly usable, and aligned with emerging European standards, serving not only public administrations but also citizens, researchers, universities, and other actors seeking reliable information about AI systems beyond the high-risk applications covered by the EU AI Act database. Although no harmonized European standard for AI transparency registers currently exists, Germany has the opportunity to establish a national model for transparent and trustworthy AI governance that is genuinely informative for the public rather than merely administrative in form.

¹⁰See e.g. <https://www.heise.de/en/news/NRW-police-modernizes-data-analysis-Who-meets-the-criteria-besides-Palantir-11273481.html>

Acknowledgements

IP, XH and MC were supported by the “TOPML: Trading Off Non-Functional Properties of Machine Learning” project funded by Carl Zeiss Foundation, grant number P2021-02-014.

References

- Ada Lovelace Institute; AI Now Institute; and Open Government Partnership. 2021. Algorithmic Accountability for the Public Sector: Learning from the First Wave of Policy Implementation. Technical report.
- Bell, A.; Nov, O.; and Stoyanovich, J. 2023. Think about the stakeholders first! Toward an algorithmic transparency playbook for regulatory compliance. *Data & Policy*, 5: e12.
- Busuioc, M.; Curtin, D.; and Almada, M. 2023. Reclaiming transparency: contesting the logics of secrecy within the AI Act. *European Law Open*, 2(1): 79–105.
- Cath, C.; and Jansen, F. 2022. Dutch Comfort: The Limits of Ai Governance Through Municipal Registers. *Techné Research in Philosophy and Technology*, 26(3): 395–412.
- Das vom BMI initiierte Beratungszentrum für Künstliche Intelligenz (BeKI). 2025. Marktplatz der KI-Möglichkeiten (MaKI): Startseite. <https://maki.beki.bund.de/a/bmi-makimo-app?kiosk>, Accessed April 2026.
- David Roth-Isigkeit. 2024. Stellungnahme im Rahmen der öffentlichen mündlichen Anhörung am 15. Mai 2024 des Digitalausschusses des Bundestags zum Thema “Nationale Spielräume bei der Umsetzung des europäischen Gesetzes über künstliche Intelligenz”. Technical report.
- Eurocities Network. 2026. Algorithm Register - Algorithmic Transparency Standard.
- European Commission. 2025. Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (AI Act). Technical report.
- European Parliament and Council. 2024. Regulation (EU) 2024/1689 (Artificial Intelligence Act). Accessed: March 2026.
- Floridi, L. 2020. Artificial Intelligence as a Public Service: Learning from Amsterdam and Helsinki. *Philosophy & Technology*, 33(4): 541–546.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; III, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Commun. ACM*, 64(12): 86–92.
- Global Partnership on Artificial Intelligence. 2024. Algorithmic Transparency in the Public Sector: A state-of-the-art report of algorithmic transparency instruments. Technical report.
- Haataja, M.; van de Fliert, L.; and Rautio, P. 2020. Public AI Registers: Realising AI transparency and civic participation in government use of AI. Technical report, Saidot.
- Kemper, J.; and Kolkman, D. 2019. Transparent to whom? No algorithmic accountability without a critical audience. *Information, Communication & Society*, 22(14): 2081–2096. [eprint: https://doi.org/10.1080/1369118X.2018.1477967](https://doi.org/10.1080/1369118X.2018.1477967).
- Leslie, M.; and Selman, C. 2024. Securing meaningful transparency of public sector use of AI: Comparative approaches across five jurisdictions. Technical report, Public Law Project.
- Lorenz, A. 2025. Konzept zur fachlichen, organisatorischen und technischen Umsetzung eines nationalen KI-Transparenzregisters. https://algorithmwatch.org/de/wp-content/uploads/2025/05/2505-AW_KI-TR-Konzept.pdf, Accessed April 2026.
- Ministerie van Binnenlandse Zaken en Koninkrijksrelaties. 2024. Algoritmeregister Publicatiestandaard.
- Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*.
- Murad, M. 2021. *Beyond the “Black Box”: Enabling Meaningful Transparency of Algorithmic Decision-Making Systems through Public Registers*. Thesis, Massachusetts Institute of Technology. Accepted: 2022-01-14T14:49:25Z.
- Nieuwenhuizen, E. 2024. Algorithm Registers: A Box-Ticking Exercise or Meaningful Tool for Transparency? *Information Polity*, 29(4).
- Open Government Partnership. 2025. Promoting Government Transparency of Algorithmic Tools in the United Kingdom.
- Soizic Pénicaut; and IA Ciudadana. 2025. Making Algorithm Registers Work for Meaningful Transparency. Technical report.
- van Vliet, M.; Schuitemaker, N.; Espana, S.; Weerd, I. v. d.; and Brinkkemper, S. 2024. Defining and Implementing Algorithm Registers: An Organizational Perspective. *ECIS 2024 Proceedings*.