The map is not the territory: conceptualising algorithm registers as transparency gateways through participatory system mapping

ANONYMOUS AUTHOR(S)*

In recent years, algorithm registers have been taken up as a means of providing a first line of transparency into algorithmic systems in public services. However, different stakeholders have different expectations of what level of transparency algorithm registers should provide. In this paper, we report on a study to better understand the role of algorithm registers as a governance instrument. To do this, we probe the municipal algorithm register of the City of Rotterdam through a case study of Avola, a decision-support tool for caseworkers' assessment of citizens' welfare benefits eligibility based on legal automation through a business rule engine. We facilitated two participatory system mapping workshops (with municipal staff and civil society organisations, N=Z) to understand to what extent the register allows stakeholders to map the algorithmic system in question. We then use these maps to inform a System-Theoretic Process Analysis (STPA) to situate the register within a wider sociotechnical governance architecture. We draw on the field of safety science to frame AI governance as a hierarchical control structure spanning technical, operational, organisational, and institutional contexts, and ask to what extent the information presented in the register can help to map actors and their relations across these contexts. We then conceptualise algorithm registers as maps that provide gateways to already existing information and mediating access for interested parties with different information needs. Reimagining the register as a system map thus alleviates the burden of completeness and instead refocuses their role as one of providing functional abstractions that are gateways to finding appropriate information sources.

CCS Concepts: • **Computing methodologies** → **Artificial intelligence**; • **Applied computing** → *Law, social and behavioral sciences.*

Additional Key Words and Phrases: policy, civil advocacy, artificial intelligence, power, justice, harms

ACM Reference Format:

Anonymous Author(s). 2025. The map is not the territory: conceptualising algorithm registers as transparency gateways through participatory system mapping. In *SIG CHI 2025 Workshop on Sociotechnical AI Governance, April 27, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 6 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

In recent years, municipal algorithm registers have become a popular instrument for providing transparency into the use of algorithmic systems in local government. Early adopters, like the Dutch city of Rotterdam, have embraced registers as a platform for informing citizens about how their municipalities use algorithms to provide services such as welfare benefits eligibility. However, recent empirical studies have shown that there are still open challenges to be resolved before these instruments can provide truly meaningful and actionable transparency. Both government auditors and academic researchers have found that these registers either disclose too little information to be useful, or are too technical for citizens to be able to understand [8, 12]. In Rotterdam, the Court of Auditors concluded from this that it is simply not clear who the register is meant to inform [12].

Manuscript submitted to ACM

Despite these shortcomings, the registers can have a 'disciplinary effect' on organisations, serving as a meaningful box-ticking exercise that motivates them to take stock of their algorithmic systems, identify their risks, and ensure that appropriate mitigation strategies are in place [8]. In this sense, registers can be understood as a meta-governance instrument which keeps an inventory of existing safeguards and governance instruments in a public-facing platform.

In this study, we ask *how can we situate algorithm registers within a wider AI governance strategy?* We seek to answer this question by engaging different stakeholder groups in a participatory system mapping exercise to probe what visibility the register provides to different actors. We use the lens of (mis)abstraction to ask what level of *information hiding* is appropriate for different actors, and to what extent *information neglect* prevents the register from facilitating meaningful accountability. We draw on safety science and systems engineering to provide a schema for sociotechnical AI governance that accounts for technical, operational, organisational, and institutional factors.

2 BACKGROUND

2.1 Transparency and the abstraction of social context

Transparency is an essentially contested concept. Its interpretation is highly variable depending on who is disclosing information and who that information is disclosed to. In a review of transparency in public administration, Meijer notes that *"transparency may contribute to accountability [...] when there are actors capable of processing the information."* [7, p.1] In the context of algorithms in the public sector, algorithm transparency can also help to foster citizens' trust in government. However, to do so it must be both accessible and interpretable [4]. Norval et al. note that if transparency is too technical, it may *"undermin[e] the disclosure's effectiveness, [..] disempower subjects, and ultimately hinder broader transparency aims"* [9].

Unfortunately, auditors and scholars alike have found that Rotterdam's algorithm register suffers from both being too technical and not technical enough, effectively not satisfying any particular audience [8, 12]. Without a clear notion of what publics the register is meant to inform, what their information needs are, and what form of transparency is meaningful for them, the register will remain unable to provide meaningful accountability. In order to better specify the ambitions of the register, it is necessary to situate it within a wider sociotechnical governance structure, attending to key actors and their relationships. To do this, we now turn to the concept of abstraction and its operationalisation in order to ground the floating signifier of transparency.

3 (MIS)ABSTRACTION

We approach the question of meaningful transparency and information disclosure through the lens of abstraction. The abstraction of social context from technical solutions can render them ineffective or even harmful, e.g. as in the development of fairness metrics for algorithmic systems [13]. However, abstractions are also a fundamental aspect of complex sociotechnical systems, where the coordination of multiple actors with often differing expertise requires abstracting away unnecessary detail in order to translate functional requirements efficiently so that others have the information necessary to do their job [1, 11]. Abstractions can help to mask unnecessary complexity in order to support efficient information exchange (*information hiding*), or to disregard nuances deemed irrelevant for a higher-level functional objective (*information neglect*) [3]. As argued in both empirical studies on algorithm registers and conceptual work on transparency, the notion of transparency in registers is currently too abstracted from the social context of the stakeholders whom that transparency is meant to inform. By excluding the information needs of the registers' users, the register performs transparency that is unintelligible to laypersons and unhelpful to expert users. The lens

of abstraction helps us to ask what level of transparency or disclosure is needed in order to provide meaningful and actionable transparency for differently situated actors.

In order to operationalise the concept of abstraction in sociotechnical systems, we draw on the analytic framework of *misabstraction* [1] to problematise the kinds of *information hiding* and *information neglect* that are enacted by algorithm registers. Anon et al [1] define a misabstraction as *a representation of an entity, phenomenon, or procedure that omits critical contextual information and renders that representation problematic when it is reintegrated into the context of the sociotechnical system for which it has been made.* We mobilise the lens of misabstraction to interrogate to what extent the information provided in the register accounts for (or omits) contextual information necessary to determine whether or not an algorithmic system has appropriate safeguards in place to prevent or mitigate emergent harms.

4 MODELING SOCIOTECHNICAL AI GOVERNANCE AS A HIERARCHICAL CONTROL STRUCTURE

We draw on Baxter & Sommerville's definition of sociotechnical systems as systems composed of social and technical components that are jointly designed for [2]. In this way, we recognise that algorithmic systems, such as those used for the determination of welfare eligibility, are constituted by an orchestration of human actors whose relations are (partly) mediated through technical artifacts, operating within specific organisational contexts, and following formal and informal rules and procedures, such as legal constraints and work practices which structure the behaviour and relations among actors within the system [10].

We can model a sociotechnical AI system and its governance apparatus as a *hierarchical control structure*. We draw on the field of safety science and systems engineering to propose an ontology for sociotechnical systems in which functional objectives (e.g. preventing harm through safety mechanisms) are achieved through information sharing and coordination among actors for the facilitation of decision-making and action-taking [1, 11]. In these arrangements, different actors (and entities) perform different actions as a means of achieving their functional objectives. Functional objectives can be specified at the macro-level of the entire system (e.g. determining welfare eligibility through the automation of the relevant legal basis) or at a micro-level specific to different sites within the broader system (e.g. a caseworker can exercise their discretionary power to overturn an automated decision-support tool when a person's dessert of welfare support may not be perfectly aligned with a strict legal interpretation). We can model information exchange, coordination, decision-making, and action-taking as a control system in which actors and entities communicate up and down the hierarchical control structure. In this hierarchical structure, control signals are sent top-down through *reference channels* in order to steer system behaviour within an operating margin of safety. These control signals are calibrated in response to bottom-up feedback signals sent through *measurement channels* which reflect the reality on-the-ground.

Through this model, we can interpret the functional objective of the algorithm register as providing visibility into the broader sociotechnical system and its governance. We now illustrate this through a case study, below.

5 CASE STUDY

We examine a case study of the municipal algorithm register of the City of Rotterdam. We study the affordances and limitations of the register by taking as an example one of the many algorithmic systems that it documents.

5.1 The Rechten Rotterdammers welfare eligibility algorithm

In order to probe the algorithm register, we take an example of one of the algorithmic systems that it documents. We consider an algorithmic decision support tool that municipal caseworkers use to determine whether or not citizens

are eligible for welfare subsidies. The system in question, Avola, is a “business rule engine" which codifies the written laws that govern welfare subsidy eligibility. The resulting law-codification implementation is validated by legal experts. Frontline caseworkers use Avola as a decision-support tool when they make their own determinations about applicants' welfare eligibility status.

6 STUDY DESIGN: PARTICIPATORY SYSTEM MAPPING

In order to *situate algorithm registers within a wider AI governance strategy*, we use Systems-Theoretic Process Analysis (STPA) [5, 6] to reconstruct the hierarchical control structure of the Avola benefits eligibility system. This preliminary system map serves as a starting point for a series of participatory system mapping workshops with municipal staff and civil society organisations. In the workshops, we will study how different actors socially situated understanding of both the Avola system enables them to reconstruct the hierarchical control structure of the Avola system. We will examine the use of abstractions in the register which serve to *hide* and *neglect* information about the system.

7 ALGORITHM REGISTERS AS GATEWAYS FOR MAPPING SOCIOTECHNICAL SYSTEMS

We conceptualise algorithm registers as maps that provide gateways to information. The task of maintaining the register thus becomes a task of coordinating among different governing task forces and mediating access for interested parties who are differently socially situated. Reimagining the register as a system map thus alleviates the burden of completeness and instead refocuses their role as one of providing functional abstractions that are gateways to unpacking more detailed information that already exists somewhere within the organisation. We suggest that seeing registers through the lens of abstraction helps to attend to the need to evaluate how well functional objectives are satisfied, and to the fact that different actors have different functional objectives.

A PRELIMINARY SYSTEM MAP

We share here a preliminary system map developed by the authors in preparation for the participatory workshops.

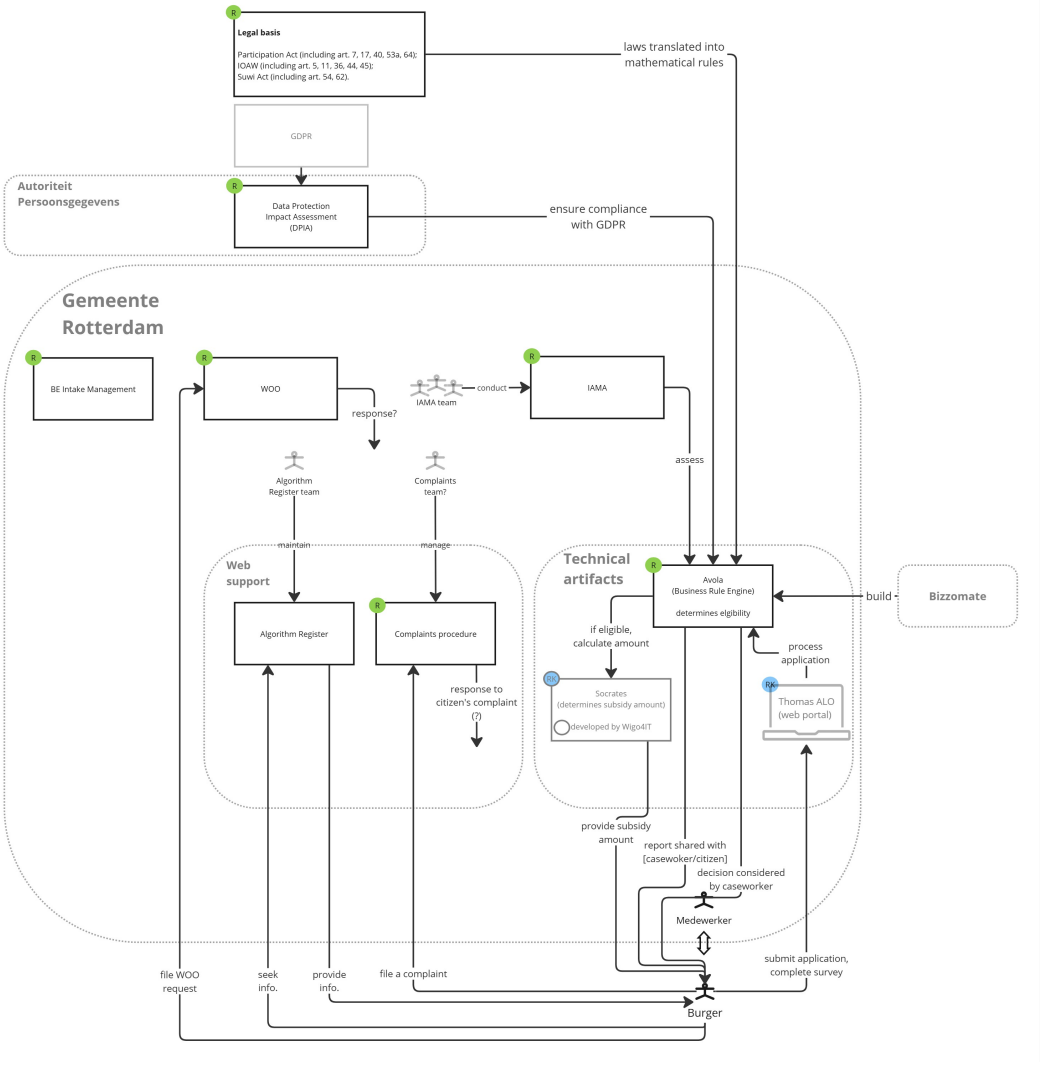

Fig. 1. The preliminary system map produced by the authors.

REFERENCES

[1] Anonymous. 2025. Sociotechnical isabstraction considered harmful. In *Under review.*

[2] Gordon Baxter and Ian Sommerville. 2011. Socio-technical systems: From design methods to systems engineering. *Interacting with computers* 23, 1 (2011), 4–17.

[3] Timothy Colburn and Gary Shute. 2007. Abstraction in computer science. *Minds and Machines* 17 (2007), 169–184.

[4] Stephan Grimmelikhuijsen. 2023. Explaining why the computer says no: Algorithmic transparency affects the perceived trustworthiness of automated decision-making. *Public Administration Review* 83, 2 (2023), 241–262.

[5] Nancy Leveson. 2004. A new accident model for engineering safer systems. *Safety Science* 42, 4 (2004), 237–270.
[6] Nancy G Leveson and John P Thomas. 2018. *STPA Handbook*. MIT Partnership for Systems Approaches to Safety and Security (PSASS).
[7] Albert Meijer. 2014. Transparency. In *The Oxford Handbook of Public Accountability*, Robert E. Goodin Mark Bovens and Thomas Schillemans (Eds.). Oxford University Press.
[8] Esther Nieuwenhuizen. 2025. Algorithm Registers: A Box-Ticking Exercise or Meaningful Tool for Transparency? *Information Polity* 0, 0 (2025).
[9] Chris Norval, Kristin Cornelius, Jennifer Cobbe, and Jatinder Singh. 2022. Disclosure by Design: Designing information disclosures to support meaningful transparency and accountability. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 679–690.
[10] Sem Nouws, Marijn Janssen, and Roel Dobbe. 2022. Dismantling digital cages: Examining design practices for public algorithmic systems. In *International Conference on Electronic Government*. Springer, 307–322.
[11] Jens Rasmussen. 1985. The role of hierarchical knowledge representation in decisionmaking and system management. *IEEE Transactions on Systems, Man, and Cybernetics* 2 (1985), 234–243.
[12] Rotterdam Court of Auditors. 2024. Kleur bekennen: vervolgonderzoek naar algoritmes. https://rekenkamer.rotterdam.nl/onderzoeken/kleur-bekennen/. Published: 2024-05-13.
[13] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. 59–68.