

Communications on Fairness Assessments in the Public Sector

An Analysis of Roles, Skills, and Responsibilities from Dutch Use-Cases

ANONYMOUS AUTHOR(S)

Algorithms used in the public sector, e.g., for allocating social benefits or predicting fraud, often require involvement from multiple stakeholders at various phases of the algorithm’s life-cycle. This paper focuses on the communication issues between diverse stakeholders that can hinder fairness assessment and potentially lead to misinterpretation and misuse of algorithmic systems. With this research objective, we conducted ethnographic research via 11 semi-structured in-depth interviews with practitioners working on algorithmic systems in the Dutch public sector, at local and national level. Applying qualitative coding analysis, we identify key elements of the communication processes that underlie fairness-related human decisions. More specifically, we analyze the division of roles and tasks, the required skills, and the challenges perceived by diverse stakeholders. Three general patterns emerge from the coding analysis. First, policymakers, civil servants, and domain experts are less involved compared to developers throughout the different phases of a system’s life-cycle. This leads to developers taking on the role of advisor and decision-maker, while they potentially miss the required skills. Second, end-users often lack the technical skills to interpret a system’s output, and rely on actors taking on developer roles for making decisions concerning fairness issues. Third, citizens are structurally absent throughout the algorithm’s life-cycle. This may lead to unbalanced fairness assessments that do not include key input from relevant stakeholders. We formalize the underlying communication processes in a conceptual framework – introducing the *phase-actor-role-task-skill* (PARTS) ontology – that can both (i) represent the communication patterns identified in the interviews, and (ii) explicitly outline missing elements in communication patterns exhibited by actors having different roles, skills, and responsibilities in a certain use case. The framework can be further extended to explore communication issues in other use cases, design potential solutions, and organize accountability with a common vocabulary.

CCS Concepts: • **Software and its engineering** → **Risk management**; *Collaboration in software development*; • **Social and professional topics** → *Computing / technology policy*.

Additional Key Words and Phrases: Communication Framework, Fairness, Transparency, Accountability, Public Sector, Qualitative User Study

ACM Reference Format:

Anonymous Author(s). 2023. Communications on Fairness Assessments in the Public Sector: An Analysis of Roles, Skills, and Responsibilities from Dutch Use-Cases. In *Boston 2023: ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, October 30 – November 1, Boston, MA, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Algorithms are increasingly being used for predicting various forms of public sector services such as allocating benefits in the domains of education, (mental) health, and detecting fraud in allowances and taxes [26, 37, 38, 51, 53]. These applications can be beneficial, but can also have detrimental consequences for citizens when used in high-stake scenarios, such as fraud detection and risk assessment. Notorious cases where incorrect predictions led to wrongful accusations of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

Manuscript submitted to ACM

53 citizen minorities are the COMPAS case in the US¹, the SyRI-case² and the Childcare Benefit Scandal in the Netherlands³,
54 which resulted in lawsuits and the latter eventually led to the resignation of the Dutch government⁴. Nowadays, the
55 problem of *fairness* in AI is widely recognized in well-established legal and ethical guidelines [10, 11, 15]. Fairness, in
56 this context, typically refers to “fair outcomes”, a principle which indicates an absence of prejudice or favoritism toward
57 an individual or group based on their inherent or acquired characteristics through algorithmic decision-making [34].
58 Most of the proposed guidelines remain generic, and the tools and methods to apply them in real-world applications
59 are missing [1, 18]. Especially, in the public sector, a disconnection exists between the proposed frameworks and the
60 additional ethical and legal frameworks that public practitioners already deal with [17]. Furthermore, data literacy at
61 public organizations may also not be mature enough to fully recognize ethical issues in data practices [43].

64 Contemporary research on fairness in AI has predominantly taken a technical approach, addressing model (dis)-
65 functionalities without taking into account the wider network of stakeholders [7, 44]. However, algorithms are always
66 the result of a process driven by (human) stakeholders’ choices and norms [48]. A solely technical approach is insufficient
67 in addressing all algorithm design choices, especially when they have an impact and involve codifying crucial social
68 values, requiring to take normative decisions [33]. Decisions about which fairness criteria are assessed and how can
69 be made throughout the algorithms’ life-cycle, from data-oriented phases (eg. data collection and labeling), to model-
70 oriented phases (eg. feature engineering, training), and execution-oriented phases (eg. deployment) [2, 25]. For example,
71 when allocating social benefits, it has to be decided which data features represent ‘eligibility’ for a social benefit [4].
72 Also, the punitive (detecting a crime) or assistive (allocating a benefit) nature of the intervention, might require choosing
73 different fairness metrics [38, 41]. Involving diverse stakeholders is also a design choice.⁵ Particularly in the public
74 sector, the need for involving diverse actors and stakeholders in algorithm development is important for ensuring that
75 public interest is prioritized, and potential harms are minimized [46].

79 In the context of algorithmic decision-making, fairness assessment is therefore understood as a process of evaluating
80 the extent to which algorithms and their outputs are free from bias, discrimination, and other forms of unfairness
81 towards different groups, individuals, or communities [4, 12, 20]. In this paper, we take a socio-technical perspective on
82 the algorithmic system and focus on the choices and practices applied at each phase of the algorithm’s life-cycle that
83 can lead to undesirable effects of bias and unfair outcomes [48]. We will analyze the internal communication processes⁶
84 involved in fairness-related human decisions, by identifying the roles, the divisions of tasks, the required skills, and the
85 potential communication challenges between diverse actors occurring throughout the algorithm’s life cycle.

88 The research questions we address are the following:

- 89 • *RQ1. Which actors, roles, tasks, and phases can be identified between multi-stakeholder interactions throughout the*
90 *algorithm’s life-cycle?*
- 91 • *RQ2. How is information about fairness assessments communicated between diverse actors involved throughout the*
92 *algorithm’s life-cycle?*

95 ¹Correctional Offender Management Profiling for Alternative Sanctions (COMPAS): the software used to predict the risk of a person recommitting a
96 crime was more inclined to falsely accuse African-American offenders than Caucasian offenders, see for instance [16, 34]

97 ²SyRI legislation in breach of European Convention on Human Rights. <https://edu.nl/xjubf>

98 ³The Dutch childcare benefit scandal, institutional racism, and algorithms. European parliamentary questions <https://edu.nl/y3h3j>.

99 ⁴Dutch Government resigns over Child Benefit Scandal <https://www.theguardian.com/world/2021/jan/15/dutch-government-resigns-over-child-benefits-scandal>

100 ⁵According to the European Commission’s Ethics guidelines on trustworthy AI [15], an important step in supporting AI fairness assessment includes
101 involving and educating all stakeholders about their roles and needs throughout the AI system’s life-cycle.

102 ⁶With internal communication we refer to communication between collaborating partners such as decision-makers, end users, developers, etc. as opposed
103 to external communication which refers to communication with the general public [17].

- RQ3. Which communication challenges can be identified for fairness assessment processes?

To answer these questions, we conducted 11 interviews with practitioners working on algorithmic systems in the Dutch public sector for fraud detection and risk assessment. We identify who makes decisions about what, and in which phase of the algorithm's life-cycle. Additionally, we also look for the (co-)occurrences of mentioned actors, roles, tasks, phases, and challenges and labeled them through *in-vivo*, descriptive, and process qualitative coding analysis [40]. From the emerging coding patterns, we identified a lack of clarity in communication among different roles and actors, a lack of feedback regarding the interpretation of the model outcome, and a lack of involvement of the relevant roles and actors at the right phase. These communication issues indicate inadequate model governance such as missing roles, skills, and information exchange which can lead to misinterpretation of model outcomes, and potentially to the inability to recognize and address fairness issues throughout the algorithm's lifecycle. Next to this descriptive effort, we build the *phase-actor-role-task-skill* (PARTS) conceptual framework to characterize the structure of the communication processes based on the coding results. The framework is provided in the form of an ontology to structure the relations between the emerging coding patterns. PARTS enables to specify the relations and interactions between the identified concepts. Moreover, PARTS is introduced with the longer-term objective to be reused and possibly extended with other use cases concerning multi-stakeholder collaborations, and to promote a common vocabulary around fairness assessments.

2 RELATED WORK

Frameworks and theories from various domains have been proposed to characterize the dynamics of interactions amongst a network of actors [32, 39]. Actor-Network-Theory (ANT) and mediation theory, for example, describe the relations and interactions within a network of (artificial and natural) actors [30–32]. Following ANT, interaction with technology is never neutral as it influences or mediates the way we carry out our tasks. On the other hand, technology is continuously mediated by our social aspects, e.g. in formulating goals and other design decisions. Another related approach to describe the context of reciprocal interactions between human actors and technology is that of the socio-technical systems (STS) [39]. STS is not reserved for technology alone, it rather stresses the interactive nature of social and technical structures within an organization and society as a whole. This term is increasingly used in the field of AI to assess fairness and ethics from a broader normative context in which actors interact and operate as opposed to focusing on individual actors alone [8, 14].

Additionally, frameworks have been proposed to investigate networks and power structures. Following the tripartite model for ethics in technology, three main roles can often be identified through their responsibilities: the *developer*, who goes about the technical input, the *user* who goes about the use of the system, and the *regulator's* role, responsible of taking the "value" decisions [36]. Research on the use of automated systems for public decision-making has shown to shift discretionary power from the regulator roles to system analysts and software designers, often making them the main decision maker [6]. Whereas decision-makers from public organizations are often involved in the procurement and deployment phase, developers, sometimes from third parties, tend to be more involved in the development phase [52]. When developers become the main decision-makers for design decisions then this can exclude those stakeholders without technical knowledge in important value decisions for the system [13, 29]. These imbalanced power dynamics can ultimately lead to a form of technocracy, where governance and (moral) decision-making are based on technological insights and may only yield technological "solutions"[36].

In the Dutch public sector, multiple stakeholders from private and public organizations often collaborate at various phases of the algorithm's life-cycle [26, 45]. Stakeholders might be involved at different moments and places depending on

157 their roles and tasks and therefore responsibilities can become diffused across a network of multiple actors simultaneously
158 [5, 52]. Related empirical field research has been conducted to investigate data practices of Dutch local governments
159 [17, 28, 43]. For example, Siffels et al. (2022) argue that with the process of decentralization, many tasks from the
160 central government were delegated to municipalities without giving them more resources and capacities. As a result,
161 municipalities use data practices to deal with additional tasks and to distribute limited (social) resources. Due to a lack
162 of data literacy, public servants are unable to recognize ethical issues and thus seek collaboration with external partners
163 which affects their ability to be transparent and responsible regarding their data projects [43]. The lack of transparency
164 and responsibility in outsourcing parts of the algorithms' life-cycle is often also referred to as the "problem of many
165 hands" [9]. In another field research, Jonk and Iren (2021) performed semi-structured interviews with key personnel
166 and decision-makers at 8 Dutch municipal organizations to investigate the actual and intended use of algorithms
167 [28]. They found that there is a lack of common terminology and algorithmic expertise not only at a technical but
168 also at a governance and operational level. The authors argue that municipalities would benefit from a governance
169 framework to guide them in the use of tools, methods, and good practices to handle potential risks. Furthermore, Fest,
170 Wieringa, and Wagner (2022) investigated how higher-level ethical and legal frameworks influence everyday practices
171 for data and algorithms used in the Dutch public sector [17]. They investigated public sector data professionals at
172 Dutch municipalities and the Netherlands Police. They found that the practicality of proposed frameworks remains a
173 challenge for practitioners because they typically do not feel competent or miss the required skill set to make decisions
174 regarding responsible and accountable data practices. Data professionals, as a result, get too much autonomy and
175 discretion in handling questions that belong at the core of public sector operations. The authors argue that efforts need
176 to be put into the operation and systematization of legal and ethical questions across the data science project life-cycle.
177 Research on public algorithms has also been conducted in the fields of Human-Computer Interaction (HCI), Science and
178 Technology Studies (STS), and Public Administration (PA) [27, 33, 42]. For instance, Saxena et al. developed a framework
179 for high-stakes algorithmic decision-making in the public sector (ADMAPS) where they qualitatively coded data from
180 in-depth ethnographic study on the daily practices of U.S. Child-welfare caseworkers and prior literature [42].

181
182
183
184
185
186
187 Lastly, tools such as ontologies⁷ can be used to characterize communication structures and knowledge exchange
188 between actors[24]. Ontologies are common tools to structure knowledge and to promote a common vocabulary within
189 a certain domain [35]. Earlier ontologies on AI risks and fairness have been proposed [19, 23]. AIRO for example is an
190 ontology that underpins the AI Act, a regulation proposed by the European Commission to tackle various forms of
191 negative impact caused by the misuse of AI [10, 23]. The Fairness metrics ontology (FMO) structures knowledge for
192 fairness notions, metrics, and the relations between them [19]. We propose that what is still missing in both ontologies
193 is a clear framework to characterize the communication (issues) between diverse actors with different skill sets that
194 underlie fairness-related human decisions.
195
196

197 We propose that what is still missing in previous work is a clear framework to characterize the communication
198 (issues) between diverse actors with different skill sets that underlie fairness-related human decisions throughout the
199 algorithm's life-cycle. Following the above frameworks and theories, we assume a perspective of a socio-technical
200 interactive network, in which fairness assessments for algorithms are understood as part of a governance structure
201 where actors with different roles interact. We conduct interviews to specifically focus on interactions by means of
202 internal communication exchange and challenges that might arise between diverse actors in the Dutch public sector.
203
204

205 ⁷An ontology can be understood from a philosophical perspective referring to "the nature and structure of reality". From a knowledge engineering
206 perspective, however, it refers to the modeling of a structure of a system — by organizing relevant concepts and relations [24]. They are also described
207 as "conceptual schemas", "formal specifications of a conceptualization" and as "the abstract and simplified view of the world we wish to represent and
208 describe in a language that is understandable by humans and/or by software agents" [3, 24].

209 The communication challenges are finally characterized in a conceptual framework in the form of an ontology. With
210 this conceptual framework, we aim to provide a structured approach to the planning and auditing of communication
211 processes around technical-decision making in the public sector.
212

213 3 METHODOLOGY

214 3.1 Semi-structured interview

215 We conducted 11 semi-structured interviews, with each one divided into three sections and lasting an hour. We
216 formulated the qualitative interview questions in an open-ended manner, where participants were able to share their
217 information in their own words whilst following a general structure of topics [21, 22]. The questions are added for
218 reference in table 4 in the appendix. The questions can be divided into three main categories:
219
220
221

- 222 1 **General**: questions about the topic of the use-case, actors involved, and the content of the respondent's (teams)
223 role. The goal of the system is identified, as well as the envisioned (end) users.
- 224 2 **Development process**: type of input, resources, tasks, and roles needed throughout the development process
225 to make informed decisions. The phases of the algorithm's lifecycle are investigated by mapping out the tasks
226 made in 5 phases: formulation, development, evaluation (go-no-go), deployment, and monitoring phase.
- 227 3 **Considerations**: perceived challenges for role and task division, potential improvements or failures of the
228 system, and communication gaps are identified. Questions about assessing errors, bias, as well as the potential
229 negative impact of the model, are asked explicitly.
230
231

232 We start by giving interviewees the opportunity to mention internal communication and fairness issues spontaneously
233 and on their own terms. The first two sections ask interviewees to describe rather generic procedures and practices
234 in the system's life-cycle. We then precisely ask about communication issues in the third part of the interview. We
235 preliminary tested the interview questions with a pilot involving colleague researchers from different disciplines. The
236 questions were deemed suitable for letting interviewees describe their process of communication and related issues.
237 The suitability of the questions was checked in terms of comprehensibility and relevance to the research questions. No
238 questions were altered afterward.
239
240

241 3.2 Case studies

242 We recruited participants who have been collaborating in multi-stakeholder projects in the Dutch public sector. We
243 specifically targeted participants working in the domains of fraud detection and risk assessment, using a repository of
244 use cases⁸ from the Dutch Ministry of interior affairs⁹ and the snowball sampling technique. These use cases were of
245 particular interest because the outcomes can directly affect citizens.¹⁰
246
247

248 For this research, ($N = 11$) interviews were conducted with ($N=10$) actors involved in Dutch social domain use cases,
249 and ($N = 1$) actor in the Dutch educational domain. Table 1 describes the participants (PP), the use case, the role they
250 filled at the time of involvement and if they have technical skills. With technical skills we refer to those that are (not)
251 educated or have (no) experience in technical science. Topics discussed for the use cases were predicting fraud amongst
252 citizens (eg. when applying for public service) ($N = 5$) or predicting the need for social benefits amongst citizens in
253
254

255 ⁸Some examples of public domain use cases in the Netherlands can also be found via the Artificial Intelligence Netherlands Coalition (NL AI Coalitie)
256 website <https://nlaic.com/use-cases> and in [26, 51].

257 ⁹Dutch Ministry of Interior Affairs and Kingdom Relations <https://www.rijksoverheid.nl/ministeries/ministerie-van-binnenlandse-zaken-en-koninkrijksrelaties>

258 ¹⁰To further confirm use-case relevance, in 2019 inspection and enforcement was identified as the biggest category for AI usage in the Dutch social
259 domain which refers to the prediction of (security) risks by identifying patterns of behavior [51].
260

Table 1. Interviewees description

Interviewee	Role	Technical background	Use case domain
PP1	Developer & Researcher	yes	Risk assessment
PP2	Program Manager	no	Risk assessment
PP3	Program Manager	no	Risk assessment
PP4	Product Owner	yes	Fraud detection
PP5	Advisor & Social researcher	no	Risk assessment
PP6	Researcher & Developer	yes	Education: face expression detection
PP7	CTO development team	no	Risk assessment
PP8	Advisor & Data science researcher	yes	Fraud detection
PP9	Developer & Researcher	yes	Fraud detection
PP10	Advisor & Researcher	no	Fraud detection
PP11	Innovation manager	yes	Fraud detection

the coming years ($N = 5$). Albeit the difference between risk and fraud cases is not always clear-cut (a fraud detection algorithm can also be used for risk and vice-versa). The categorization of the use cases is therefore based on the punitive (fraud detection) or assistive nature (risk assessment) nature of the policy measure. Before conducting the interviews, participants received some example questions and a short description of the research. At the start of the interview, participants gave their consent for collaboration.

3.3 Qualitative coding analysis

To provide insights into the role and task division amongst actors in multi-stakeholder projects, we performed qualitative coding analysis to label key *codes* from the interview output¹¹. We use in vivo¹², descriptive¹³ and process coding¹⁴ to identify the process of communication exchange between diverse actors, as well as the practices and choices made at each stage of the algorithm's life-cycle. Beforehand, both analysts agreed that particular attention should be drawn to describing the role of the stakeholders, depending on the task, phase, and on potential challenges that might occur. For example, if a participant were to mention that "person X is a developer and performs bias analysis in the development phase" we label the actor, the role, the task, and the phase. We also pay attention to the direction of information exchange between actors (eg. person X hands over model outcomes to person Y). Furthermore, communication challenges and information that was repeated or pointed out by participants as essential for fairness assessment were coded. For the sake of supporting a common terminology, we took a broad range for the formulation of the codes. For example, participants mentions for engineers, coders, and data scientists are all considered as the role of developers. Also, for

¹¹ Saldaña (2013) describes "A code in qualitative inquiry is most often a word or short phrase that symbolically assigns a summative, salient, essence-capturing, and/or evocative attribute for a portion of language-based or visual data." In addition, a code can be understood as a researcher-generated construct that symbolizes the construct and assigns an interpreted meaning.

¹² In Vivo coding is also named "literal coding" and refers to a word or short phrase from the actual language found in the qualitative data record eg. terms used by participants themselves [40, 47].

¹³ descriptive coding refers to summarizing the basic topic of a passage of qualitative coding in a word (noun) or short phrase [40].

¹⁴ process coding refers to "action coding" which implies action from more simple observable activity (eg. reading) to more general conceptual action (such as adapting) [40].

313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364

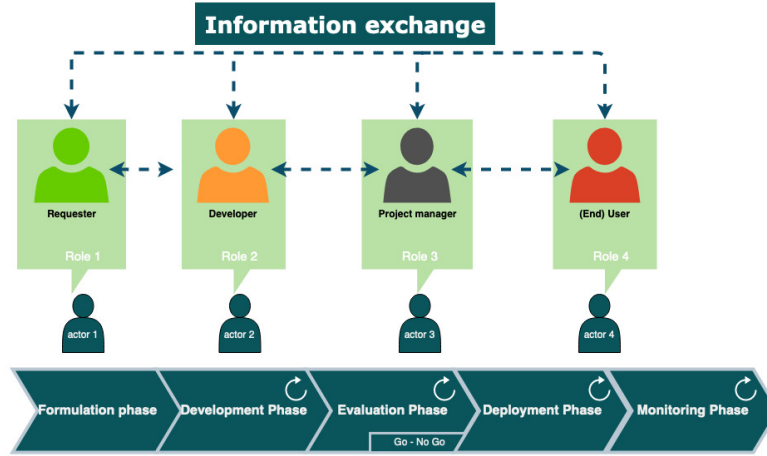


Fig. 1. Key phases of a system’s life-cycle, and their actors and roles, identified from the interview transcripts.

the phases, we considered the testing phase, modeling, and experimentation phase as the development phase. These are grouped because they were not purposely separated but used interchangeably by participants in the interviews. In figure 1, an overview can be found of examples for the identified key phases, actors and roles from the interview transcripts.

The coding analysis is done in multiple cycles, where in each round of coding pieces of text are merged or split into categories. Two of the authors perform a separate coding analysis, to reduce the impact of personal bias. One performed the coding analysis by hand whilst the other used a coding analysis tool¹⁵. Each analyst then identified the roles, tasks, phases, and challenges from the interview transcript independently. We finally compared both coding analyses to identify discrepancies or alignments, as described in the Results section. Mentions of coding are done by counting the (co-)occurrence for roles and phases, roles and tasks, and roles and challenges.

3.4 Constructing a Conceptual framework

Based on the codes identified with the qualitative coding analysis from the interview transcripts, we organize the main elements in a conceptual framework for the general communication processes between diverse actors that underlie fairness-related decisions. We build the conceptual framework using the method of constructing an ontology. The framework is therefore built in an iterative manner presuming that it can be revised and extended for future work when more input is available. We follow the work by Noy & McGuinness (2001)¹⁶ on creating ontologies [35] and take other ontologies and conceptual frameworks as a reference [19, 23, 50]. The ontology promotes a common terminology so that diverse actors can contribute to creating transparency and accountability throughout their system’s development practices.

A Creating concepts from codes We use the main identified codes from the qualitative coding analysis to make concepts, also known as classes, in our ontology. As mentioned in the previous sections, we focus on identifying roles, actors, tasks, phases, and challenges. We purposely separate actors from roles to see how many roles

¹⁵Atlas.ti: The Qualitative Data Analysis & Research Software <https://atlasti.com/>

¹⁶"Ontology development 101: A Guide to Creating Your First Ontology"

actors can take on for a certain task. These concepts are the main focus of the ontology for representing the communication structure.

B Defining Descriptions We add descriptions to each concept to agree on common definitions. Descriptions are based on answers we got from the interview transcripts, the definitions we found from documents provided by the European Commission on Trustworthy AI, and from other sources in the literature [11, 15, 23, 36, 49, 50].

C Defining properties and relations between concepts We describe the internal structure of the concepts by adding properties. For instance, for actor, we use a property to describe their public or private affiliation. We determine the relations between the concepts to describe the generic process of communication. For example, an actor always takes on a certain role and is involved during a certain phase, when performing a certain task. Since we identified various challenges from the interview transcripts, we decide that challenges can occur for all the other defined concepts.

D Competency Questions We determine the scope of the ontology by thinking of competency questions that a knowledge base derived from the ontology should be able to answer. We use our research questions as guidance to answer and formulate competency questions for who (actors, roles) does what (tasks) when (phase), and which information exchange is needed for that.

E Creating instances: We create generic instances of concepts for the main concepts. By doing this step, we aim to describe the concepts in relation to each other. This is done by taking a phrase from the interview transcript eg. a phrase that contains a challenge, and filling in concepts for specific slot values. We go back to the other steps if the instance calls for revision of the ontology.

4 RESULTS

4.1 Semi-structured interview

For the majority of use cases ($N = 10$), the procurement for the algorithm came from government organizations and municipalities. Furthermore, in all use cases, there were multiple stakeholders involved with varying technical expertise—from social workers to developers, researchers, program managers, and third parties. Finally, in most use cases ($N = 10$), the envisioned end-users of the systems were policymakers or other domain experts at municipalities with minimal or no technical expertise. There was one exception for the educational domain, which envisioned teachers as end-users. 6 interviewees mentioned having a technical background themselves, but not necessarily in AI or software development.

4.2 Qualitative coding analysis

4.2.1 Qualifying the stakeholders. Figures 2, 3, and 4 illustrate which roles were mentioned the most based on co-occurrence per Phase, Task, and Challenge. Overall, the developer role was mentioned the most ($N = 189$), followed by End-users ($N = 107$) and Policy Makers ($N = 92$). More formulation of roles, phases, tasks, and challenges were identified but mentioned less and thus not shown here. Figure 2 shows that developers are most prominent in the development, evaluation, and formulation phase and less in the deployment and monitoring phases. PP1, who filled a developer role reported “we don’t monitor what the municipalities are doing with the results.” and “feedback is needed on how the results will be used in deployment”. End-user and policy-makers are reportedly the same in almost all use cases. They are also the second highest in occurrences for phases. Moreover, figure 2 demonstrates that the monitoring phase $N = 10$ is mentioned the least throughout the interviews whereas the evaluation phase is mentioned the most $N = 90$.




Fig. 2. Mentions of Roles (relative) per Phases. The total amount of occurrences per phase is mentioned on the y-axis (from least to most). Note that the developer role (blue) is mentioned most in all phases, followed by end-users(green) and policy-makers(red) and subjects (orange) least.

The subject role occurs the least for phase and task and is in almost all cases filled by citizens. Participants reported the need for more citizen involvement and being more transparent to citizens throughout the phases algorithms' life-cycle. PP2, a program manager, mentioned that *"the plan is to check eg. with research labs how citizens can give feedback on the model. But currently, nothing is envisaged"*. On involving citizens, another researcher (PP10) working on fraud detection, reported that *"it depends on the type of AI. If it has an impact on citizens or uses a lot of data from citizens, it would be relevant to include a focus group of citizens from the beginning but it is less relevant for road repairs."* Other roles are also mentioned that need to be more involved throughout the phases. PP9, a developer, reported that *"For the future, we could incorporate stakeholders at earlier stages in the development to see what the potential sources of bias are."*

Additionally, requester roles are not mentioned frequently and are mostly involved in the formulation phase for funding or initiating the project. A requester often resides at external governmental organizations such as ministries and collaborates with municipalities. The subject role as well as the requester role were never described as end-users.

In figure 3 it is shown that the developer role is mentioned most for tasks, e.g. technical decision-making, researching, bias analysis, and consulting, and the least for model usage. This indicates that actors taking on developer roles are the most prominent in making decisions throughout the algorithm's life-cycle. On what type of decisions one makes, a PP9 who filled a developer role at the time, says that they *"decided on how to improve accuracy and handling issues. For instance, gathering more data diverse to handle bias"*. Developer PP1 mentioned that they *"define and chose metrics for the models."* and that these *"are defined in collaboration with the municipality but choosing metrics and trimming down after input was decided by the two of their team."*

Program managers and product owners (PM/PO) often work in the same research teams as developers and are either hired externally or internally by a public requester. They are the main decision-makers for involving (external) stakeholders such as advisors, requesters, and end-users, and sometimes being the only ones in direct contact with them to answer questions. PM/PO's are often reported to supervise developers in technical decision-making but rely on their judgment for bias, fairness, and risk factor analysis. Program manager (PP2) said that for handling error rates and biases they *"rely on the technical teams' judgment"*. Also they mentioned that *"the technical colleagues give advice when the model is good enough, but it's a bit of a grey area. We also rely on literature"*. PP3 mentioned that they *"don't go*

469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520

Main Role occurrence per Task




Fig. 3. Mentions of Roles (relative) per Task. The total amount of occurrences per task is mentioned on the y-axis (from most to least). Note that technical decision-making is mentioned the most. Moreover, the developer role (blue) is mentioned most for all tasks except for model usage.

to users” for error or bias analysis but “talk to the data scientist”. Product owner (PP4) said that “it is time intensive to explain [bias analysis] to stakeholder users. Bias analysis is sometimes so complex, even as an expert I sometimes don’t understand it, and it takes a lot of time”.

Advisors are often mentioned to consult on the system in the evaluation phase before deployment, or when the project is halted. Advisors are presented as giving advice on 1) domain knowledge, 2) technical knowledge, or 3) ethical knowledge. Advisor roles are sometimes combined with developer roles. A third-party developer can be hired to analyze the code, give technical advice or even build the model. PP3 added that they “hired an external bureau for auditing and investigating the algorithm”. Also because they “could not get reliable predictions because the social domain changes all the time, and it’s hard to keep track of these changes—for example in social support—and how that impacts the system”. PP10 mentioned that “an external company was hired to develop the model for the municipality”, which made the “data ecosystem quite complex”. A social domain expert and researcher at the municipality (PP5) mentioned that they “were involved to give feedback as an involved bystander. But it was hard for someone like me to understand what the difference between implementation and design is and what that means for real-life implications”.

In figure 4 the relative occurrences for roles per challenge are demonstrated. Most communication challenges were reported amongst end-users/policymakers and developers due to lack of interpretation, missing involvement, and (risk) oversight in a certain phase. Program manager (PP3) mentioned that is a challenge that “we don’t know if governments

521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572

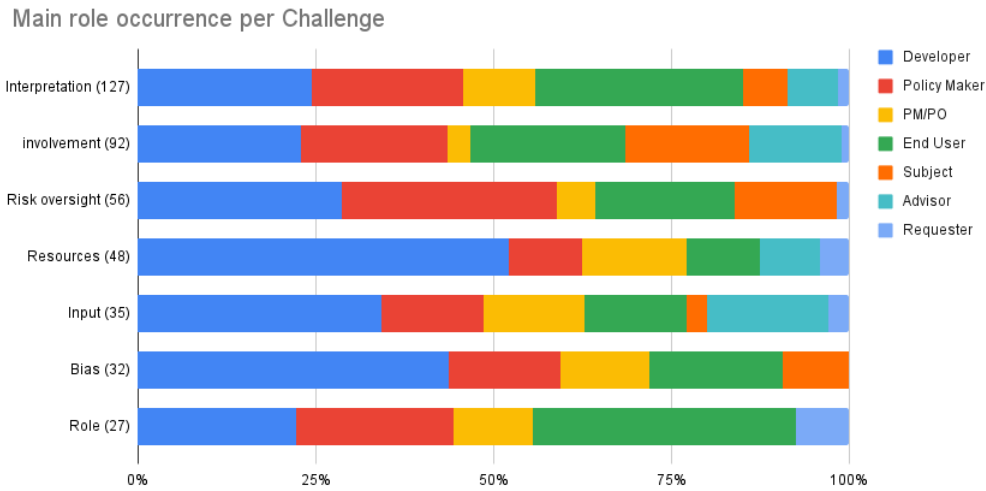


Fig. 4. Mentions of Roles (relative) per Challenge. The total amount of occurrences per challenge is mentioned on the y-axis (from most to least). Note that interpretation and involvement issues were mentioned most. Moreover, developer (blue) is mentioned most followed by end-users (green) and policy-makers (red).

and municipalities can understand the model”. Also, PP1 mentioned with regard to model usage at the municipality that “it’s hard to get a focused answer on how they are going to use the model and what the results will be.” Therefore, more input is needed on the interpretation and use of the results from end-users in the deployment phase. Vice versa on feature selection and bias analysis in the development phase end-users require input from developers. On the matter of involvement, another program manager (PP2) reported that “More frequent and streamlined collaboration with the municipality is needed” and that they “would like to have closer contact with the municipality.” PP1 also confirmed that in their case “the municipality is too loosely involved in the project, more involvement is needed” when asking about challenges regarding role division.

End-users and policy-makers are often reported to miss the technical skills to understand the uncertainty of predictions and limitations of the model in real-world settings. For example, PP10 mentioned that the most difficult challenge is the “gap between data scientists and policymakers. How to make sure that what is developed is being well understood and useful for those of non-tech background.” PP1 also confirmed that the “Most important risk is that the model will not be used or is misinterpreted. For example, mixing up correlation and causality might lead to not helping people at risk of poverty.” PP6 mentioned that “People could trust the model blindly and mistake it for a decision-making tool”, and lastly, PP8: “Not sure if the inspectors fully understood why certain cases were flagged as misuses or put on the list.”

With regard to biased decisions, another participant (PP9) mentioned that “there should be more focus on asking users what policymakers perceive as risks and biases” and that it is “difficult for them to understand that there are many different interpretations. What it really means to be a ‘true positive’, is this person really a fraud, or was this person not able to fill in the forms properly?”. Misinterpretation and misuse of the system by end-users and policy-makers were mentioned most in the monitoring and deployment phases. PP6 confirms that “training for users is needed, to remind users not to rely on

573 *the tool but that the decision is up to them.*” This also explains why policy-makers and end-users are mentioned a lot
574 throughout all phases because of their co-occurrence with challenges in certain phases.

575 Program managers and product owners are also hesitant to involve citizens for a lack of appropriate frameworks, and
576 potentially hostile attitudes. “*There is a long history with the citizen council for consultation and it is usually conflict-based.*
577 *It’s hard to make fruitful collaboration, getting them to understand the issues and getting them out of anger mode.*” said a
578 product owner (PP4) working at the municipality. PP7, on previous involvement of a citizen council: “*they said no on the*
579 *feasibility of the model from the municipality. They did not get it. It was more of a general no to technology instead of*
580 *asking a targeted question*”.

583
584 **4.2.2 Qualifying the communication processes.** Our qualitative coding analysis shows the importance of differentiating
585 individual stakeholders by their *roles* and *skills*, meaning, having specific skills to match particular roles diverse actors
586 need to have within the process. Fairness assessments are impacted when individual *actors* with specific *roles* or *skills*
587 are missing in the communication framework.

588 For example, *actors* with technical *skills* often assume the *role* of policy-makers. Yet, they may miss the required
589 *skills* in the domain of application. For instance, the developers of a fraud detection system may decide on error metrics,
590 bias issues, and test sets to measure them. Yet, they may not have the expertise in demographics, finance, and forensics
591 to ensure that their technical decisions are fair or allow them to assess fairness outcomes appropriately. On the other
592 hand, the policy-maker role may not have the technical skills to understand bias or error metrics. In this case, we can
593 consider that *actors* with the *role* of policy-makers are missing. Also, *actors* with practical *skills* in the domain of fraud
594 are missing.

597 As we reflected on our qualitative coding results, specific patterns of issues in communication frameworks emerged.
598 We were able to identify key elements of communication frameworks that are particularly critical to assessing fairness.
599 We found that issues arise depending on how these elements are combined with other elements, to form communication
600 frameworks with similar patterns and structures. These communication patterns are not easily described with qualitative
601 coding analysis only. By solely counting the codes for (co-) occurrences, we could not capture the type of relation
602 between codes that are important for determining the process of communication. For instance, a communication
603 challenge could be that a role is missing in a certain phase, which would also count as a co-occurrence for a role within
604 a phase. Their structures have variations that are not easy to identify in the narratives of interview transcripts. The
605 communication patterns are also not easy to document in written form only. Therefore, we turned to ontology modeling
606 for documenting the communication patterns we identified. In the next section, we provide a more structural analysis
607 of the communication patterns using the PARTS ontology. PARTS was specifically designed for—but not restricted to—
608 modeling the results of our qualitative coding analysis.

613 **4.3 The PARTS ontology**

615 The Phase-Actor-Role-Task-Skills (PARTS) ontology models the structure of communication frameworks and the basic
616 elements for the communication issues identified in our qualitative coding analysis. It comprises rather generic concepts,
617 such as *Phase*, *Actor* and *Information Exchange*, and concepts that are more specific to fairness assessment, such as *Role*,
618 *Task*, and *Skill*.

620 The communication frameworks require at least 2 generic elements to materialize the communication process:
621 the stakeholders who exchange information (represented with the concept *Actor*), and the act of communicating
622 (represented with the concept *Information Exchange*). We found that the relevant context of the communications can be
623

625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676

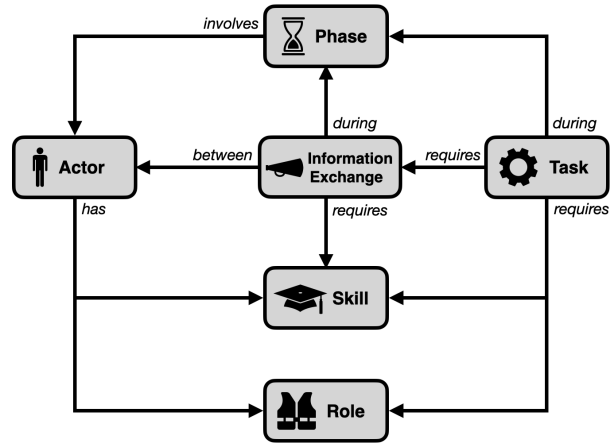


Fig. 5. Basic concepts of the PARTS ontology for characterizing the structure of communication frameworks, and the issues we identified. For example, fairness assessment Tasks may be missing at certain Phases of a system’s life-cycle. Or Actors may not have the right Skills or Role when performing fairness assessments.

Concept	Description	Relation	Property
Task	Actions carried out by Role	Requires Role, Skill, Info exchange, during Phase	
Role	Function or a part filled by Actor	Actor Has Role, Task requires Role	
Actor	Entities that perform Tasks, actively or passively	has Role, Skill, Phase involves Actor, Info Exchange between Actor	affiliation: public or private, is missing: yes/no
Phase	indicates the evolution of the system from conception through retirement	Involves Actor, Task during Phase, Info exchange during Phase	
Skill	Professional ability, expertise, or knowledge needed in practice to successfully complete a specific task	Info exchange requires Skill, Task requires Skill, Actor has Skill	is missing: yes/no
Info exchange	Communication transfer between actors	requires Skill, between Actor, during Phase, Task requires Info exchange	is missing: yes/no

Table 2. Concept definitions

modeled with 4 additional concepts (Phase, Task, Role, Skill). We found that these simple concepts were sufficient for modeling generic communication challenges that we observed in several interviews. The ontology can be extended with additional concepts, e.g., to represent the affiliation of Actors, or the information that is exchanged. We decided that further research is needed to further refine the ontology, and that adding more concepts at this stage comes at the risk of making the ontology less applicable or generalizable to new contexts.

We thus selected these 6 concepts from the codes of our qualitative analysis of the interview transcripts. Table ?? provides the definitions and properties of these concepts, and Figure 5 shows their relationships. Figure 6 provides an example of the instances of the ontology that we derived from the interviews.

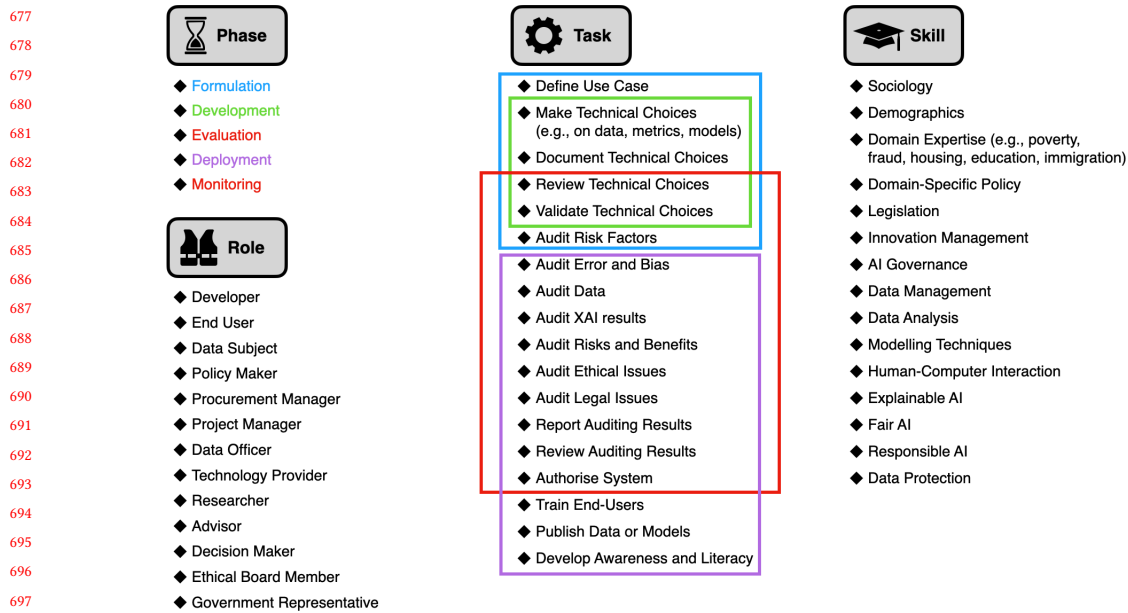


Fig. 6. Generic instances of the PARTS ontology, identified in several interviews. The generic instances of *Tasks* can occur during different *Phases*, e.g., as the colored boxes show. The *Tasks* occurring during the *Phases* of ◆ *Evaluation* and ◆ *Monitoring* can largely overlap. Instances of *Actors* are not generic, they represent private individual persons (e.g., the interviewees themselves).

Role types	Description
Developer	research, design, and/or develop algorithms
Policy-maker	responsible for designing and overseeing the carrying out of policy and social decisions
Project manager/ Product owners (PM/PO)	supervise the projects for the development of the system and oversee documentation checks and balances
End-user	(in)directly engage with the system and use algorithms within their business processes to offer products and services to others
Subject	organization or entity that is impacted by the system, service, or product
Advisor	give constructive feedback on the system throughout the life-cycle
Requestor	who are the main client and investor for the use-case

Table 3. Main roles description

We found that a few simple properties and relations between concepts were sufficient to represent the communication issues we observed in the interviews. The property *requires* assigned to the fairness assessment *Tasks* allows representation of the *Roles*, *Skills* and *Information Exchange* needed to perform the task. This part of the communication framework is essential to define the high-level requirements for accountability and transparency. The practical implementation of these requirements can be modeled by associating *Actors* with specific *Skills* to the *Roles* they take in fairness assessment *Tasks*.

729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780

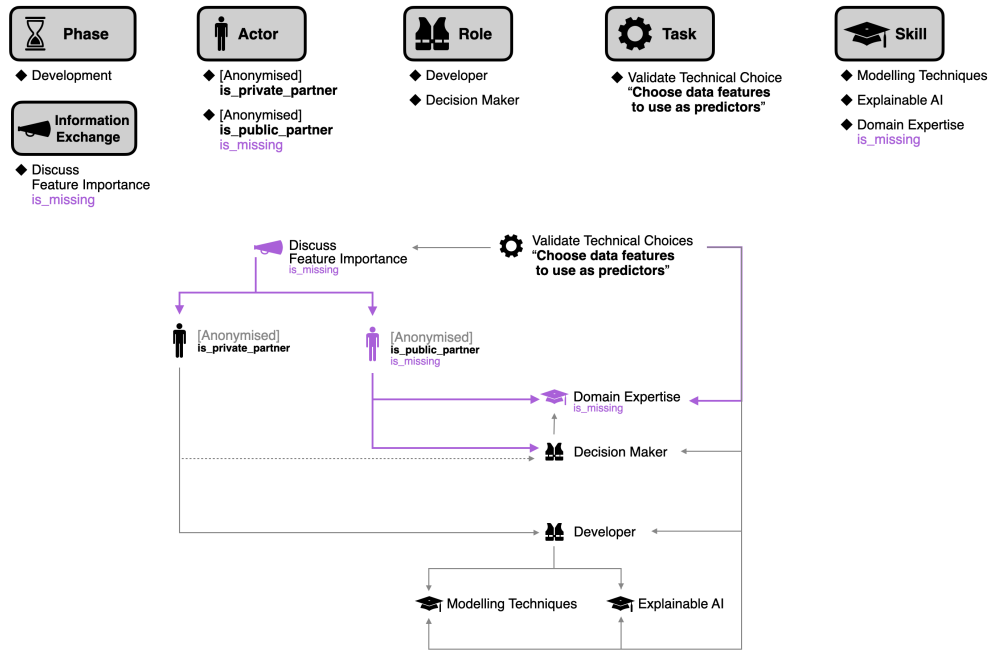


Fig. 7. Example of a generic communication challenge observed in interviews, with elements of the communication framework (top) and their structure (bottom). Existing elements are in black, and missing elements are in pink. For the *Task* of ♦ *Validating technical choices*, the *Role* of ♦ *Decision Maker* is missing an *Actor* who has *Skills* in the ♦ *Domain Expertise*, and who is independent of the ♦ *Developer*.

With the properties *during* and *involves*, the ontology can associate *Tasks*, *Information Exchanges* and *Actors* to the *Phases* of the system’s life-cycle. This provides a generic temporal overview of the communication and assessment processes, e.g., to reflect on the assessment *Tasks* or *Actors* (e.g., stakeholders from specific institutions) that may be missing at specific *Phases*.

Our main finding is that adding the property *is_missing* to any element of the communication framework (i.e., to any of the 6 concepts) is of great interest for expressing the issues that interviewees mentioned. We represented the elements of communication frameworks using these 6 concepts precisely because issues arise if any of these types of elements are missing. For example, an *Actor* may be missing specific *Skills*, or a fairness assessment *Task* may be entirely missing. Communication issues also arise from missing a *Information Exchange*, as shown in Figure 7.

The concepts we selected, and their descriptions, are in correspondence with those provided by the European Commission on trustworthy AI, and other sources in the literature [11, 15, 23, 36, 49, 50]. For example, the developers and end-users (as instances of *Actor*) are, respectively, those who “research, design, and/or develop algorithms” and those who “(in)directly engage with the system and use algorithms within their business processes to offer products and services to others” [15]. Data subjects (also instances of *Actor*) are “an organization or entity that is impacted by the system, service or product” [23].

5 DISCUSSION

The patterns identified from the coding analysis indicate a heavy reliance on the (technical) skills of the developer role in the development, formulation, and evaluation phases. This means that developers likely take on most tasks and decisions throughout the life-cycle which may potentially influence fairness outcomes.

These results are based on the co-occurrence of mentions provided by participants in interview transcripts. However, participants tend to forget involved internal and external stakeholders in the development process or don't see some as influential for choices, practices, and protocols in the development process. This does not mean that forgetting a certain role or actor fully reflects the actual governance structure or experienced communication challenges. Participants may not have oversight, may not be willing to provide specific details, or were perhaps steered by the interview questions.

Besides, we recognize that in practice, formulations for roles and groups can vary and can be diffused. For instance, in some cases, actors identified as developers would primarily identify themselves as a researcher – who also carry out developer tasks or vice-versa. Moreover, counting (co-)occurrences is not enough to assess the nature and structure of interactions within a network. As mentioned in the results, sometimes an occurrence would be counted for role and phase when actually the content of the code would say that 'role X was missing in phase Y' and thus was more an occurrence for the challenge than for phase. Therefore, the PARTS ontology was needed to provide an overview and to characterize the interaction between roles and the challenges within a bigger network.

The patterns additionally indicate a need for increased input of end-users, policy-makers, and developers at the right phase, in particular during design choices and output interpretation. Even when there is involvement, participants indicate that the necessary skills are missing to provide for informed decisions. There is uncertainty about whether the algorithm actually provides what it is supposed to perform, coupled with weak support and guidance for users, limited supervision, advice, and a lack of feedback loop for AI developers. These findings demonstrate a growing information gap between technical experts on one side and domain experts and policy-makers on the other side. Presenting algorithms and features in an understandable way to end-users and policy-makers is increasingly needed, aimed specifically at a clear identification of potential risks, failures, and errors of the AI model.

Moreover, within patterns, there is no clear governance or process structure regarding roles, actors, information flows, and responsibilities during the algorithm's life-cycle. The findings revealed it was not always clear who makes final (mostly policy) decisions on the further development or use of algorithms, or what is the (legal, procedural, information) basis for such decisions. This seems to be due to the lack of actors filling the right roles at the right phase; which leads to actors taking on multiple roles at once they are not fully equipped for, consequently leading to a discretionary imbalance. For example, developers tend to take on multiple roles such as advisor, researcher, and decision-maker, without being fully aware of the business logic and domain expertise. In addition, policy-maker or social worker roles tend to coincide with end-user roles without being allocated a clear mandate for making informed value decisions. These findings are not new and confirm previous work done on the growing discretionary power of developers and unclear role divisions for automated decision tool usage in policy-making settings eg. [6, 36].

The issues of information gaps, unclear role divisions, and co-occurrence were characterized by the PARTS ontology. PARTS is a communication tool that can explicitly structure and outline missing elements such as missing roles, skills, and involvement. Furthermore, it promotes a common terminology beyond solely the technical perspective of fairness assessments but also on the matter of accountability. We purposefully kept the ontology as simple as possible because every public sector use case will have its own (normative) context, specific governance, and communication structure.

Also, we do not claim that fairness issues are resolved with PARTS. We solely emphasize that governance structures, roles, and communication processes between diverse actors are of importance when addressing fairness assessments.

This paper shows that when responsibilities and roles are unclear, and there is a lack of communication or miscommunication, this can lead to misinterpretation of the system output and even to misuse of the algorithmic system. Finally, the coding analysis, as well as the generic patterns presented in PARTS are based on ($N = 11$) interviews. More data is needed in the future to verify and specify these findings. Future work would require increased development and use of practical frameworks for communication, role and responsibility distribution, and governance processes.

6 CONCLUSION

This paper examines the roles of different stakeholders engaged in the process that leads to algorithm procurement and development in the Dutch public sector. In particular, we focus on the analysis of communication patterns exhibited in fairness assessments. The field research we conducted points to inadequate or missing explicit guidance or rules regarding roles, skills, and information exchange throughout the algorithm's life-cycle which leads to misinterpretation of model outcomes and potentially to unfair outcomes. The results have been characterized using the Phase-Actor-Role-Task-Skill ontology and can be used to indicate a lack of involvement and feedback between developer, end-user, and policy-maker roles. End-users and policy-makers often lack the technical skills to interpret the system's output or to estimate potential fairness issues. Therefore, they rely on the technical skills of developers for making apparent technical decisions such as bias and fairness assessments, potentially influencing policy outcomes. As a result, developers take on extra roles such as advisor and decision-maker while potentially lacking the business logic and domain expertise required. This also results in developers being the most prominent in most tasks and phases of the algorithm's lifecycle. Lastly, we observe that citizens are structurally absent throughout the algorithmic life-cycle, even though it is mentioned that their involvement is needed in the future for balanced fairness assessment. The PARTS ontology is built to describe these patterns and to be used and extended by domain experts: to audit and design frameworks to deal with communication issues, to support transparency and accountability amongst diverse actors, and to promote the creation of a common vocabulary on fairness assessments.

REFERENCES

- [1] Kasun Amarasinghe, Kit T. Rodolfa, Hemank Lamba, and Rayid Ghani. 2020. Explainable Machine Learning for Public Policy: Use Cases, Gaps, and Research Directions. *CoRR* abs/2010.14374 (2020). arXiv:2010.14374 <https://arxiv.org/abs/2010.14374>
- [2] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. 2019. Software engineering for machine learning: A case study. In *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 291–300.
- [3] Grigoris Antoniou, Enrico Franconi, and Frank Van Harmelen. 2005. Introduction to semantic web ontology languages. *Reasoning web* 3564 (2005), 1–21.
- [4] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org. <http://www.fairmlbook.org>.
- [5] Mark Bovens. 2007. 182 Public Accountability. In *The Oxford Handbook of Public Management*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199226443.003.0009> arXiv:https://academic.oup.com/book/0/chapter/292909174/chapter-ag-pdf/44514193/book_34524_section_292909174.ag.pdf
- [6] Mark Bovens and Stavros Zouridis. 2002. From Street-Level to System-Level Bureaucracies: How Information and Communication Technology is Transforming Administrative Discretion and Constitutional Control. *Public Administration Review* 62, 2 (2002), 174–184. <https://doi.org/10.1111/0033-3352.00168> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/0033-3352.00168>
- [7] Shea Brown, Jovana Davidovic, and Ali Hasan. 2021. The algorithm audit: Scoring the algorithms that score us. *Big Data & Society* 8, 1 (2021), 2053951720983865.
- [8] Amit K. Chopra and Munindar P. Singh. 2018. Sociotechnical Systems and Ethics in the Large. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society* (New Orleans, LA, USA) (*AIES '18*). Association for Computing Machinery, New York, NY, USA, 48–53. <https://doi.org/10.1145/3209838.3209848>

- 885 1145/3278721.3278740
- 886 [9] Jennifer Cobbe, Michael Veale, and Jatinder Singh. 2023. Understanding aschccountability in algorithmic supply chains. *arXiv preprint arXiv:2304.14749* (2023).
- 887
- 888 [10] European Commission. 2021. Proposal for a regulation of the European parliament and of the council: laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206>
- 889
- 890 [11] European Commission, Content Directorate-General for Communications Networks, and Technology. 2020. The Assessment List for Trustworthy Artificial Intelligence (ALTAD) for self assessment. <https://doi.org/10.2759/002360>
- 891
- 892 [12] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*. 797–806.
- 893
- 894 [13] J. Danaher. 2016. The threat of algocracy: Reality, resistance and accommodation. *Philosophy & Technology* 29, 3 (2016), 245–268.
- 895 [14] Mateusz Dolata, Stefan Feuerriegel, and Gerhard Schwabe. 2022. A sociotechnical view of algorithmic fairness. *Information Systems Journal* 32, 4 (2022), 754–818.
- 896
- 897 [15] Content European Commission, Directorate-General for Communications Networks and Technology. 2019. Ethics Guidelines for Trustworthy Artificial Intelligence. <https://doi.org/10.2759/346720>
- 898
- 899 [16] Tracy L Fass, Kirk Heilbrun, David DeMatteo, and Ralph Fretz. 2008. The LSI-R and the COMPAS: Validation data on two risk-needs tools. *Criminal Justice and Behavior* 35, 9 (2008), 1095–1108.
- 900
- 901 [17] Isabelle Fest, Maranke Wieringa, and Ben Wagner. 2022. Paper vs. practice: How legal and ethical frameworks influence public sector data professionals in the Netherlands. *Patterns* 3, 10 (2022), 100604.
- 902
- 903 [18] Luciano Floridi, Josh Cows, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, and More Authors. 2018. AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines: journal for artificial intelligence, philosophy and cognitive sciences* 28, 4 (2018), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- 904
- 905 [19] Jade S Franklin, Karan Bhanot, Mohamed Ghalwash, Kristin P Bennett, Jamie McCusker, and Deborah L McGuinness. 2022. An Ontology for Fairness Metrics. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 265–275.
- 906
- 907 [20] Sorelle A Friedler, Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P Hamilton, and Derek Roth. 2019. A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the conference on fairness, accountability, and transparency*. 329–338.
- 908
- 909 [21] L. A. Fujii. 2018. *Interviewing in Social Science Research, A Relational Approach*. Routledge.
- 910 [22] De Goede, Bosma, and Pallister-Wilkins. 2019. *Secrecy and Methods in Security Research A Guide to Qualitative Fieldwork*. Routledge.
- 911 [23] Golpayegani, Harshvardhan, and Lewis. 2022. *AIRO: An Ontology for Representing AI Risks Based on the Proposed EU AI Act and ISO Risk Management Standards*. ResearchGate.
- 912 [24] Nicola Guarino, Daniel Oberle, and Steffen Staab. 2009. What is an ontology? *Handbook on ontologies* (2009), 1–17.
- 913 [25] Mark Haakman, Luis Cruz, Hennie Huijgens, and Arie van Deursen. 2020. AI lifecycle models need to be revised. an exploratory study in fintech. *arXiv preprint arXiv:2010.02716* (2020).
- 914
- 915 [26] Hoekstra, Chideock, and Van Veenstra. 2021. TNO Rapportage Quicksan AI in the Publieke sector II. <https://www.rijksoverheid.nl/documenten/rapporten/2021/05/20/quicksan-ai-in-publieke-dienstverlening-ii>
- 916
- 917 [27] Naja Holten Møller, Irina Shklovski, and Thomas T Hildebrandt. 2020. Shifting concepts of value: Designing algorithmic decision-support systems for public services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–12.
- 918
- 919 [28] Eric Jonk and Deniz Iren. 2021. Governance and Communication of Algorithmic Decision Making: A Case Study on Public Sector. In *2021 IEEE 23rd Conference on Business Informatics (CBI)*, Vol. 1. IEEE, 151–160.
- 920
- 921 [29] Pratyusha Kalluri. 2020. Don't ask if artificial intelligence is good or fair, ask how it shifts power. *Nature* 583, 7815 (July 2020), 169–169. <https://doi.org/10.1038/d41586-020-02003-2>
- 922
- 923 [30] Bruno Latour. 1992. Where are the missing masses? The sociology of a few mundane artifacts. *Shaping technology/building society: Studies in sociotechnical change* 1 (1992), 225–258.
- 924
- 925 [31] Bruno Latour. 1994. On technical mediation. *Common knowledge* 3, 2 (1994).
- 926 [32] Bruno Latour. 1999. On recalling ANT. *The sociological review* 47, 1_suppl (1999), 15–25.
- 927 [33] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–35.
- 928
- 929 [34] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- 930
- 931 [35] N. Noy and B. McGuinness. 2001. Ontology development 101: A guide to creating your first ontology. https://protege.stanford.edu/publications/ontology_development/ontology101.pdf
- 932
- 933 [36] van de Poel and Royakkers. 2011. *Ethics, technology, and engineering : an introduction*. Wiley-Blackwell, United States.
- 934 [37] Kit T Rodolfa, Hemank Lamba, and Rayid Ghani. 2021. Empirical observation of negligible fairness-accuracy trade-offs in machine learning for public policy. *Nature Machine Intelligence* 3, 10 (2021), 896–904.
- 935
- 936

- 937 [38] Kit T. Rodolfa, Erika Salomon, Lauren Haynes, Iván Higuera Mendieta, Jamie Larson, and Rayid Ghani. 2020. Case study: predictive fairness to
938 reduce misdemeanor recidivism through social service interventions. In *FAT* '20: Conference on Fairness, Accountability, and Transparency, Barcelona,*
939 *Spain, January 27-30, 2020*, Mireille Hildebrandt, Carlos Castillo, L. Elisa Celis, Salvatore Ruggieri, Linnet Taylor, and Gabriela Zanfir-Fortuna (Eds.).
940 ACM, 142–153. <https://doi.org/10.1145/3351095.3372863>
- 941 [39] Günter Ropohl. 1999. Philosophy of socio-technical systems. *Society for Philosophy and Technology Quarterly Electronic Journal* 4, 3 (1999), 186–194.
- 942 [40] J. Saldaña. 2013. *The coding manual for qualitative researchers*. SAGE.
- 943 [41] Pedro Saleiro, Benedict Kuester, Loren Hinkson, Jesse London, Abby Stevens, Ari Anisfeld, Kit T. Rodolfa, and Rayid Ghani. 2018. Aequitas: A Bias
944 and Fairness Audit Toolkit. <https://doi.org/10.48550/ARXIV.1811.05577>
- 945 [42] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2021. A framework of high-stakes algorithmic decision-making for
946 the public sector developed through a case study of child-welfare. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–41.
- 947 [43] Lotje Siffels, David van den Berg, Mirko Tobias Schäfer, and Iris Muis. 2022. Public Values and Technological Change: Mapping how municipalities
948 grapple with data ethics. *New Perspectives in Critical Data Studies* (2022), 243.
- 949 [44] Giovanni Sileno, Alexander Boer, Geoff Gordon, and Bernhard Rieder. 2021. Like circles in the water: Responsibility as a system-level function. In
950 *AI Approaches to the Complexity of Legal Systems XI-XII: AICOL International Workshops 2018 and 2020: AICOL-XI@ JURIX 2018, AICOL-XII@ JURIX*
951 *2020, XAILA@ JURIX 2020, Revised Selected Papers XII*. Springer, 198–211.
- 952 [45] J. Spierings and S. van der Waal. 2020. Algoritme: de mens in de machine - Casusonderzoek naar de toepasbaarheid van richtlijnen voor algoritmen.
953 https://waag.org/sites/waag/files/2020-05/Casusonderzoek_Richtlijnen_Algoritme_de_mens_in_de_machine.pdf
- 954 [46] Logan Stapleton, Devansh Saxena, Anna Kawakami, Tonya Nguyen, Asbjørn Ammitzbøll Flügge, Motahhare Eslami, Naja Holten Møller, Min Kyung
955 Lee, Shion Guha, Kenneth Holstein, et al. 2022. Who Has an Interest in “Public Interest Technology”? Critical Questions for Working with Local
956 Governments & Impacted Communities. In *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social*
957 *Computing*. 282–286.
- 958 [47] Anselm L Strauss. 1987. *Qualitative analysis for social scientists*. Cambridge university press.
- 959 [48] Harini Suresh and John V. Guttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *EAAMO*
960 *2021: ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, Virtual Event, USA, October 5 - 9, 2021*. ACM, 17:1–17:9.
961 <https://doi.org/10.1145/3465416.3483305>
- 962 [49] Damian A Tamburri, Willem-Jan Van Den Heuvel, and Martin Garriga. 2020. Dataops for societal intelligence: a data pipeline for labor market skills
963 extraction and matching. In *2020 IEEE 21st International Conference on Information Reuse and Integration for Data Science (IRI)*. IEEE, 391–394.
- 964 [50] Willem van Hage, Véronique Malaisé, Roxane Segers, Laura Hollink, and Guus Schreiber. 2011. Design and use of the Simple Event Model (SEM).
965 *Web Semantics: Science, Services and Agents on the World Wide Web* 9 (2011), 128–136. <https://doi.org/10.1093/oxfordhb/9780199226443.003.0009>
- 966 [51] A.F.E. Van Veenstra, S. Djafari, F. Grommé, B. Kotterink, and R.F.W. Baartmans. 2019. Quickscan AI in the Publieke dienstverlening. <http://resolver.tudelft.nl/uuid:be7417ac-7829-454c-9eb8-687d89c92dce>
- 967 [52] Maranke Wieringa. 2020. What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability. In
968 *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (Barcelona, Spain) (FAT* '20)*. Association for Computing Machinery,
969 New York, NY, USA, 1–18. <https://doi.org/10.1145/3351095.3372833>
- 970 [53] Ben Williamson. 2016. Digital education governance: data visualization, predictive analytics, and ‘real-time’ policy instruments. *Journal of education*
971 *policy* 31, 2 (2016), 123–141.

Table 4. Questionnaire

989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040

Questions	Notes
Institution / Department What is your (team's) role? Who do you work with (directly)? Domain and topic of use case Start and end date What type of system is being developed for the use case? What is the goal of the system? External partners developing technology for the use case (if any) Who are the (end) users - are they directly involved in the development process?	Name of entity/department Description team/staff involved brief Intended use/aim
II. Development process Could you guide us through the process of development by mapping out phases - and specific actions in each phase? Could you guide us through the decisions made about the system and by whom? What kind of decisions do you and your team make? (could you give an example) What input is needed /do you use to make decisions as a reference point What kinds of exchanges are needed in your decision process? Are there other teams or (external) stakeholders involved in the decision-making process? How do you support the decision process of your collaborators with your output?	e.g. involvement in deciding on: the goal of the system design of the system (metrics/labeling/test/training/error) evaluation of system monitoring deploying system e.g. handbook, training, expert group?
III. Considerations How can the development process be improved for the following, from your perspective? What are the most difficult challenges and risks of failures for the system? How are these challenges and risks measured assessed and monitored? What information is needed (by whom) to handle these challenges and risks? Who is consulted for this information? What is your role in the process of addressing challenges and risks? Could any issues occur that might halt the development process, (if so could you give an example of how are these go/no-go decisions determined?) In real-life applications, could there be specific risks or negative impacts for individuals or social groups? Is error analysis / Bias analysis performed for negative impacts (if so how is this done and could you give an example)? Once the algorithm is deployed in practice, what kind of human oversight is available to control for error, bias, or negative impacts? What procedures and recourses, if any, are available for addressing the negative impacts of the system? Do you have access to explanations or training on the risks for individuals and social groups, e.g., from your colleagues or from external ²⁰ experts?	eg. information exchange, role division, handling error rates and biases, handling risks, responsibilities
IV. Follow-up questions Who is in charge / responsible for mitigating measures on respecting privacy and data protection? For instance, is there a valid legal basis for processing personal data? Are there cybersecurity or privacy-preserving measures deployed to preserve privacy and data security? If no challenges (or very few) concerns are mentioned in 4, provide a scenario?	e.g. complaints about the output; security breaches; what if the training set is not representative, high error rates