

# What does it mean to be a responsible AI practitioner: An ontology of roles and skills

Shalaleh Rismani  
McGill University  
Canada

AJung Moon  
McGill University  
Canada

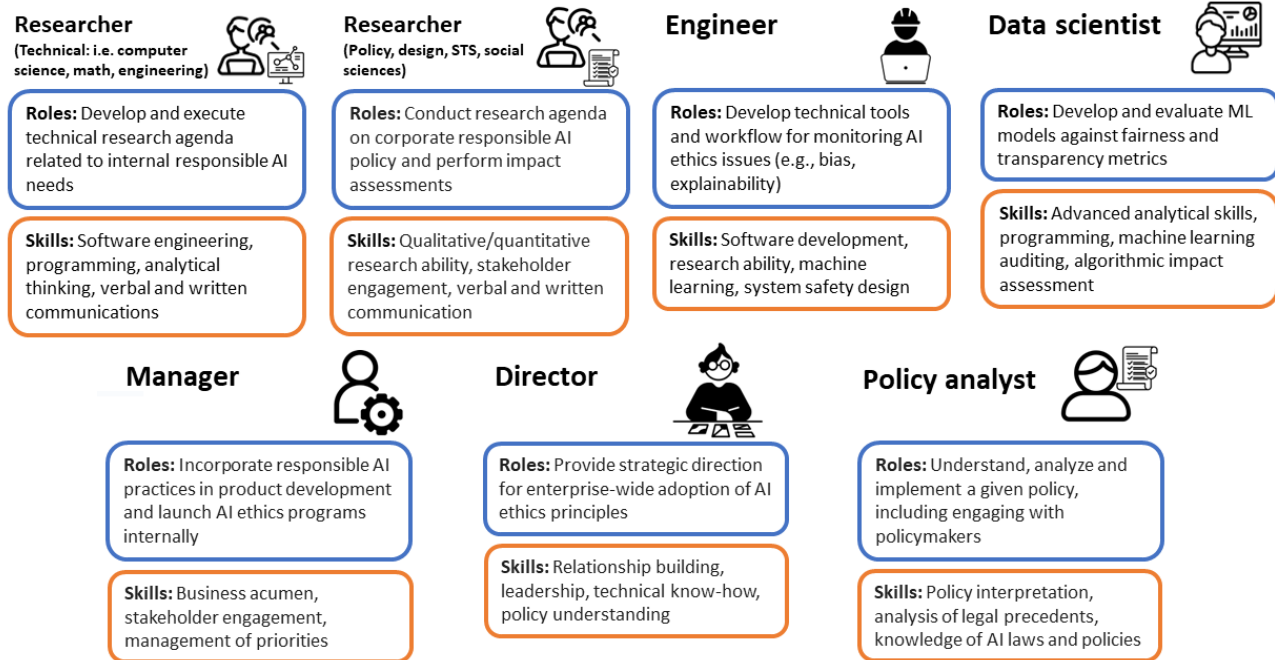


Figure 1: Existing roles and skills expected of responsible AI practitioners (AI ethicists)

## ABSTRACT

With the growing need to regulate AI systems across a wide variety of application domains, a new set of occupations has emerged in the industry. The so-called responsible Artificial Intelligence (AI) practitioners or AI ethicists are generally tasked with interpreting and operationalizing best practices for ethical and safe design of AI systems. Due to the nascent nature of these roles, however, it is unclear to future employers and aspiring AI ethicists what specific function these roles serve and what skills are necessary to serve the functions. Without clarity on these, we cannot train future AI ethicists with meaningful learning objectives.

In this work, we examine what responsible AI practitioners do in the industry and what skills they employ on the job. We propose an ontology of existing roles alongside skills and competencies

that serve each role. We created this ontology by examining the job postings for such roles over a two-year period (2020-2022) and conducting expert interviews with fourteen individuals who currently hold such a role in the industry. Our ontology contributes to business leaders looking to build responsible AI teams and provides educators with a set of competencies that an AI ethics curriculum can prioritize.

## CCS CONCEPTS

• **Social and professional topics** → **Computing profession.**

## KEYWORDS

Competency Framework, Responsible AI Practitioner, Education

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

AIES '23, August 08–10, 2023, Montréal, QC, Canada

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0231-0/23/08.

<https://doi.org/10.1145/3600211.3604702>

## ACM Reference Format:

Shalaleh Rismani and AJung Moon. 2023. What does it mean to be a responsible AI practitioner: An ontology of roles and skills. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '23)*, August 08–10, 2023, Montréal, QC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3600211.3604702>

## 1 INTRODUCTION

With the rapid growth of the AI industry, the need for AI and AI ethics expertise has also grown. Companies and governmental organizations are paying more attention to the impact AI can have on our society and how AI systems should be designed and deployed responsibly [23, 31, 42]. From 2015 onward, a series of AI ethics principles [31], in-depth auditing toolkits [11, 39, 46], checklists [5, 35], codebases [4, 8], standards and regulations [1, 6] have been proposed by many different international actors. Several communities of research and practice such as FATE (Fairness, Accountability, Transparency, and Ethics), responsible AI, AI ethics, AI safety and AI alignment have emerged. This general movement towards responsible development of AI has created new roles in the industry referred to as *responsible AI practitioners* in this paper. The primary mandate of these roles is understanding, analyzing, and addressing ethical and social implications of AI systems within the business context. The emergence of these roles challenges technology companies to curate these roles and teams. Leaders in AI-related organizations need to identify, recruit and train appropriate candidates for such roles. As the demand to fill such roles continue to increase, educators need effective means to train talent with the right set of skills.

Recently, scholars examined the common roles responsible AI practitioners serve [25, 55], explored the challenges that they face [40, 47], and criticized the problematic nature of the accountability mechanisms that relate to these roles [19]. Moreover, others highlight the myriad practical challenges facing the development of a comprehensive training program to fill such roles [14, 26, 45]. However, there is a lack of empirical research investigating the types of roles, corresponding responsibilities, and qualifications that responsible AI practitioners have in the industry. To address these gaps, we examine the following research questions:

- **RQ1:** What are the types of roles and responsibilities that responsible AI practitioners hold in the industry?
- **RQ2:** What are the skills, qualifications, and interpersonal qualities necessary for holding such roles?

We address these questions by conducting a two-part qualitative study. We examined 79 job postings from March 2020 to March 2022 and conducted expert interviews with 14 practitioners who currently hold these roles in the industry. Learning from fields of competency-based recruitment and curriculum development, we propose an ontology of different occupations and an accompanying list of competencies for those occupations.

As illustrated in Figure 1, our ontology outlines seven occupations that responsible AI practitioners hold in the industry: researcher (of two kinds), data scientist, engineer, director/executive, manager, and policy analyst. For each occupation, the ontology includes a list of responsibilities, skills, knowledge, attitudes, and qualifications. We find that while the roles and responsibilities held by responsible AI practitioners are wide-ranging, they all have interdisciplinary backgrounds and are individuals who thrive in working with individuals from different disciplines. We discuss how educators and employers can use this competency framework to develop new curricula/programs and adequately recruit for the rapidly changing field of responsible AI development.

## 2 BACKGROUND

With the increased media reporting and regulation requirements around social and ethical issues of AI-based products and services [7, 38, 48, 51, 54, 56], the role of a responsible AI practitioner has emerged as a demanding position in the technology industry. In this section, we provide an overview of debates about these roles and existing educational programs that aim to train future responsible AI practitioners. We discuss how existing competency frameworks treat the role of a responsible AI practitioner and highlight the gaps we address in this work.

### 2.1 Emergence of the responsible AI practitioners

Considering the nascency of AI ethics as a domain, only a few scholars have characterized occupations held by responsible AI practitioners [36, 57]. For instance, Gambelin frames the role of an AI ethicist as "an individual with a robust knowledge of ethics" who has the responsibility and the ability to "apply such abstract concepts (i.e. ethical theories) to concrete situations" for the AI system. According to Gambelin, an AI ethicist in the industry also needs to be aware of existing policy work, have experience in business management, and possess excellent communication skills [25]. Gambelin identifies bravery as the most important characteristic of an AI ethicist as they often need to "shoulder responsibility" for potential negative impacts of AI in the absence of regulation.

Moss and Metcalf investigated practices and challenges of responsible AI practitioners in Silicon Valley and described them as "ethics owners" who are responsible for "handling challenging ethical dilemmas with tools of tech management and translating public pressure into new corporate practices" [40]. Echoing Moss and Metcalf's seminal work on examining AI industry practices, a growing body of empirical work highlights that responsible AI practitioners face challenges such as misalignment of incentives, nascent organizational cultures, shortage of internal skills and capability, and the complexity of AI ethics issues when trying to do their day-to-day tasks [41, 47, 48, 50, 55]. Furthermore, only large technology companies often have the necessary resources to hire responsible AI practitioners [52]. Small and medium-sized companies struggle to access such expertise and rely on openly available information or hire external consultants/auditors as needed [19, 52]. This has given rise to AI ethics as consulting and auditing service [10, 18, 34].

While challenges in operationalizing responsible AI practices are an active area of research, there is a gap in understanding the role and necessary competencies of responsible AI practitioners in the industry.

### 2.2 Qualifications to be a responsible AI practitioner

The emergence of auditors in the field of responsible AI emphasizes the need for formal training and certification of such roles in the industry [19]. This raises a few practical questions: Who is qualified to take these roles? How should these individuals be trained? Are existing computer science, engineering, and social science curricula prepare individuals for such roles?

Educators responded to this need by developing a range of educational programs and curricula [14, 24, 28, 44, 58]. In a survey of

the curricula for university courses focused on AI ethics, Garrett et al. emphasize that such topics should be formally integrated into the learning objectives of current and new courses [26]. On the other hand, as Peterson et al. describe, discussing social and ethical issues in computer science courses remains a challenge [43]. They propose pedagogues for fostering the emotional engagement of students in the classroom as a solution [43].

Recognizing the importance of interdisciplinary approaches in AI ethics, Raji et al. argue that computer science is currently valued significantly over liberal arts even in the research area of fairness of machine learning systems [45]. Furthermore, they state that the perceived superiority culture in computer science and engineering has created a "new figure of a socio-technical expert", titled "Ethics Unicorns" - full stack developers, who can solve challenging problems of integrating technology in society.

This overemphasis on computer science expertise and the trend toward integrating ethics content in existing technical curricula may be problematic if these efforts do not match the skills and disciplinary needs of the industry. It raises questions about whether the educational backgrounds of responsible AI practitioners today are indeed in computer science. In this work, we inform the curriculum development efforts across a diverse range of disciplinary areas by understanding these roles in the industry and outlining the attributes, qualifications, and skills necessary for holding them.

### 2.3 Competency frameworks in AI and AI ethics

Competency frameworks are useful tools for human resource management (i.e. recruitment, performance improvement) and educational development (i.e. new training programs and curriculum development in universities) [2, 53]. Competency frameworks highlight different competencies required for a profession and link these competencies to skills and knowledge. According to Diana Kramer "competencies are skills, knowledge and behaviours that individuals need to possess to be successful today and in the future" [49]. This definition frames our discussion of competency in this paper.

Competency frameworks help governmental and non-governmental organizations keep track of the type of skills their employees/general public need in the short and long term. Educators use these frameworks to update existing curricula and develop appropriate learning objectives. On the other hand, business leaders and human resource professionals use these frameworks for their recruitment practices.

Today's existing competency frameworks do not sufficiently represent roles and competencies of a responsible AI practitioner. For example, O\*NET is United State's national program for collecting and distributing information about occupations [9]. O\*NET-SOC is a taxonomy that defines 923 occupations and they are linked to a list of competencies. Searching the taxonomy for "ethics", "machine learning", "data", "security", and "privacy" leads to minimal results such as "information security analysis", "data scientist and "database architect". The dataset do not include occupation titles such as machine learning engineer/researcher or data/AI ethics manager.

ESCO, the European skills, competencies, qualifications, and occupation is the European and multilingual equivalent of US's O\*NET [21]. ESCO contains 3008 occupations and 13890 skills. Searching for the above terms leads to more relevant results such as computer vision engineer, ICT intelligent system designer, policy

manager, corporate social responsibility manager, ethics hacker, data protection officer, chief data officer, and ICT security manager. However, emerging occupations relevant to AI and AI ethics have not been well-represented in these established, Western competency frameworks.

As a response, a number of new AI competency frameworks have recently been developed. One such enabler is the series of projects funded by the Pôle montréalais d'enseignement supérieur en intelligence artificielle (PIA), a multi-institutional initiative in Montreal, Canada aimed to align educational programs with the needs of the AI industry. Six projects related to AI competency frameworks were funded – including the work presented in this paper. This resulted in an overarching AI competency for postsecondary education that includes ethical competencies [13], and a competency framework specific to AI ethics skills training [16]. Bruneault et al., in particular, created a list of AI ethics competencies based on interviews of university instructors/professors already teaching courses related to AI ethics across North America.

Our work complements these collective efforts by providing a framework that represents the needs of the industry expressed in recent AI ethics-related job postings and the realities of the jobs AI ethics practitioners hold in nonprofit and for-profit corporations today.

## 3 METHODOLOGY

Practitioners and scholars of different domains typically create competencies frameworks using a process most appropriate for their needs. However, many follow a version of the process highlighted by Sanghi [49]. The steps of the process are: 1) Define the purpose and performance objective of a position, 2) Identify the competencies and behaviors that predict and describe superior performance in the job, 3) Validate selected competencies, 4) Implement/integrate competencies and 5) Update competencies.

In this work, we focus on answering questions raised in the first two steps about the objectives of *responsible AI practitioner* roles and skills/qualities required to perform well in these positions. We take a two-pronged approach to understand the nature of emerging roles under the broad category of *responsible AI practitioners* in the industry. Firstly, we reviewed and analyzed job postings related to our working definition of *responsible AI practitioner*. Secondly, we interviewed individuals who are responsible AI practitioners in the industry today. We then synthesized data collected from these two sources through thematic analyses. We present our proposed competency framework in Section 4. This study was approved by the Research Ethics Board of our academic institution.

### 3.1 AI Ethicist Job Postings Review

We collected and analyzed 94 publicly available job postings over the period of March 2020 to March 2022. The job postings included a range of job titles, including researcher, manager, and analyst. The following sections describe the process for collecting, selecting, and analyzing these job postings that led to the development of the ontology of responsible AI practitioner roles and skills.

**3.1.1 Collection of job postings.** To collect "AI ethicist" job postings, we searched and scraped three job-finding websites, including LinkedIn, indeed.com, and SimplyHired, every two months from

March 2020 to March 2022. We used the following search terms: AI ethics lead, Responsible AI lead, AI ethics researcher, data OR AI ethicist and fairness OR transparency researcher/engineer. Considering that search results only showed a few relevant job postings, we also collected job postings that came through referrals, including mailing lists such as FATML, 80000hours.org, and roboticsworldwide.

After scanning all the resulting job postings with the inclusion criteria, we gathered a total of 79 job postings for thematic analysis. We included the job postings that were published within our data collection period, were situated in the industry (including not-for-profit organizations), and outlined responsibilities with regards to implementing AI ethics practices in a given sector.<sup>1</sup>

**3.1.2 Analysis.** Using Braun and Clarke's thematic analysis methodology [15], we analyzed the job postings with the coding scheme illustrated in Table 1. The lead author created this coding scheme after reviewing all the postings. The coding scheme was also informed by frequently used categories across competency frameworks explained earlier in section 2.3.

The codes were generally split into four key elements: the company environment, responsibilities in the given occupation, qualifications, and skills. The codes of "company environment" and "qualifications - interdisciplinarity" are unique to this coding scheme due to their prevalence in the postings' content.

After developing the first draft of the coding scheme, a student researcher was trained to use this scheme and coded 10% of the job postings. The student researcher's analysis using the coding scheme was consistent with the lead researcher's analysis of the same set of job postings. The discussion between the lead and student researcher helped clarify the description and examples for each code. However, there were no new codes that were added to the scheme. The lead author updated the coding scheme and coded the entire set of postings using the new scheme.

**Table 1: Coding scheme for Job Posting Analysis**

Code
<b>Company environment</b>
<b>Occupation</b>
occupation - non-technical roles
occupation - technical roles
occupation - title
<b>Qualifications</b>
qualifications - education
qualifications - experience
qualifications - interdisciplinarity
<b>Skills/competency</b>
skills/competency - attitudes/values
skills/competency - knowledge
skills/competency - language skills
skills/competency - skills

<sup>1</sup>The table outlining the inclusion and exclusion criteria is in the supplemental material.

## 3.2 Expert interviews

The job postings provide a high-level analysis of the required skills and competencies expressed by recruiters; however, they may not represent the reality of these roles. Therefore, we conducted 14 interviews with experts who currently hold responsible AI practitioner positions in the industry. The focus of the interviews was on understanding the responsibilities, qualifications, and skills necessary for these roles. Considering the objective of this research project on the type of roles and skills, we did not acquire any demographic information about the participants in these roles. This also ensured that we can maintain the anonymity of these participants considering that a limited number of people hold these positions.

**3.2.1 Recruitment.** We compiled a list of potential interview candidates through (a) referrals within the authors' professional network and (b) we used similar search terms as the ones highlighted for job postings to look for people who currently hold these positions. Moreover, we also considered people from the industry who had accepted papers at relevant conferences such as FAccT and AIES in 2020 and 2021. The suitable participants:

- worked for a minimum of three months in their role;
- held this position in the industry or worked mainly with industrial partners;
- held managerial, researcher, technical positions that are focused on implementing responsible AI practices within the industry.

We did not interview researchers or professors in academic institutes and only interviewed those holding positions at nonprofit and for-profit companies. While we only used the search terms in English to find interview participants for practical reasons, we did not limit our recruitment efforts to a geographical region given the limited number of individuals holding these roles across the industry. We recruited and conducted interviews from June 2021 - February 2022.

**3.2.2 Interview protocol.** The primary researcher conducted all fourteen interviews. All of the interviews were 45 to 60 minutes in length. The interviewer first described the project and obtained the participant's consent. The interview was semi-structured with ten questions focused on exploring the following four topics:<sup>2</sup>

- Background and current role
- Situation your work, projects in AI ethics
- Skills, knowledge, values
- Looking into the future

**3.2.3 Data Analysis.** We recognize that this research reflects our positionality and biases as academics in North America. Furthermore, the data we collected were all in English and they were representative of job postings and positions in companies situated in North America and Europe. We were not able to collect data on job postings and candidates representing existing efforts in Asia and the Global South. Furthermore, we recognize that the roles in this field are continually shifting. Therefore, this ontology is only a snapshot of the roles and skills that responsible AI practitioners have and are recruited for today. Further iterations on these types of frameworks will be necessary in the future as these roles evolve.

<sup>2</sup>The detailed interview protocol is included in the supplementary material.

Finally, this study focuses on examining responsibilities, qualifications, and skills required of today's practitioners independent of their demographic factors (e.g., gender, age). We recognize the importance of representing a demographically diverse group of individuals and their experiences in qualitative research such as ours. Once responsible AI practitioners become a common occupation held by many, future studies should include demographic factors as part of similar investigations.

### 3.3 Author reflexivity and limitations

We recognize that this research reflects our positionality and biases as academics in North America. Furthermore, the data we collected were all in English and they were representative of job postings and positions in companies situated in North America and Europe. We were not able to collect data on job postings and candidates representing existing efforts in Asia and the Global South. Furthermore, we recognize that the roles in this field are continually shifting and see this ontology as only a starting point for understanding the roles that responsible AI practitioners take and their necessary skills. We emphasize the need to have further iterations on these types of frameworks. Finally, this study focuses on examining responsibilities, qualifications, and skills required for such roles independent of the demographics of individuals who are currently holding these roles. We recognize the importance of having a demographically diverse group of individuals across many occupations and suggest that future studies also examine issues around demographic diversity in roles related to responsible AI.

## 4 PROPOSED COMPETENCY FRAMEWORK FOR RESPONSIBLE AI PRACTITIONERS

From our analysis, we developed a preliminary competency framework that captures seven classes of existing occupational roles and several emerging classes of occupations. Figures 2 and 3 show how each occupation type was represented in the job postings and interviews. Three of the occupations require technical expertise (researcher, data scientist, and engineer), two require policy expertise (researcher, policy analyst), and the remaining two are managerial (manager, director). In the following sections, we provide a detailed description of the responsibilities, skills, qualifications, and qualities for each of these roles.

### 4.1 Researcher (technical)

The most common class of occupations found in the job postings was that of a researcher focused on technical aspects of fairness, explainability, safety, alignment, privacy and auditability of AI systems (24 job postings, 2 interviews). Employers represented in this dataset were looking to hire researchers at varying levels of seniority (assistant, associate and principal). The main responsibilities of these researchers are split into four main categories: conducting research, communicating their findings, working with other teams (internally and externally), and developing novel solutions for identified problems. As expected, research directions set by these researchers need to support company-specific needs, and there is an emphasis on communication between researchers and product, legal and executive teams.




Figure 2: Distribution of occupations represented in the job postings dataset




Figure 3: Distribution of occupations represented in the interviews

**Skills.** The researchers in this group need to have a mix of technical skills (i.e. software engineering and programming languages such as Python), research skills (i.e. analytical thinking and synthesis of complex ideas), and leadership skills (i.e. leading and guiding fellow researchers). The dataset from the job postings emphasized equally all these skills, and more senior positions emphasized leadership skills. A senior researcher explained that they look for "different research skills" depending on the project; however, they generally look for "some background in machine learning, statistics, computer science or something of that nature" and hire candidates that have some "interdisciplinary background". The data from the postings and the interviews show a strong emphasis on good verbal and written communication skills. Participants highlighted the ability to publish in academic venues and some emphasized the ability to communicate with different audiences internally (i.e. product teams and executives) and externally (policy-makers and executives). A technical researcher emphasized the importance of

"convincing stakeholders" and creating "strategic collaborations" by communicating with practitioners with "diverse" backgrounds.

**Qualifications.** The job postings mainly aim to attract candidates who have a PhD in computer science or a related field. Few of the job postings accept a master's in these fields, whereas some do not highlight a specific degree and mainly focus on necessary skills and knowledge. The majority of postings have a heavy emphasis on the required experience. Interview participants also emphasized the importance of experience. A research manager expressed that they are not necessarily looking for a "PhD in computer science". They are looking for candidates with experience in "leading and executing a research agenda", working with different people and teams, synthesizing and "communicating challenging concepts", and practicing software engineering. Some postings highlight experience with implementing AI ethics-related concepts. However, this was often listed as a preferred qualification rather than a required one. Similarly, researchers we interviewed, echoed the importance and value of having a publication record in "Fairness, Accountability, Transparency, and Ethics (FATE) communities" such as ACM Conference on Fairness, Accountability, and Transparency (FAccT) and AAAI/ACM Conference on AI, Ethics, and Society (AIES).

**Interpersonal Qualities.** The most common attitude/value was the aptitude and interest to *collaborate and work in an interdisciplinary environment*. A researcher emphasized that the current conversations are "engineering focused" and they actively incorporate perspectives from *social science and philosophy by collaborating with experts in these areas*. The most desired value was "curiosity to learn about [responsible AI] problems". Many of the participants highlight other values and attitudes such as "passion" towards building safe and ethical AI systems, *willingness to manage uncertainty and challenges, creativity, and resourcefulness*.

## 4.2 Data scientist

The data scientist occupation is represented in 10 job postings in our dataset, and none in the interviews. The job postings seek to fill traditional data scientist roles with an added focus on examining responsible AI-related issues. The common responsibilities outlined for these positions are a) to collect and pre-process data, and b) to develop, analyze, and test models – these are typical of existing data science roles. However, the job postings emphasize the position's responsibility to test machine learning models for AI ethics concerns such as fairness and transparency. Data scientists who work in the responsible AI domain have additional non-conventional roles. These roles include understanding and interpreting existing regulations, policies, and standards on the impact of AI systems and testing the systems' capability for elements covered in these policies. They also need to work with technical and non-technical stakeholders to communicate findings, build capacity around responsible AI concepts and engage them as needed.

**Skills.** The job postings put a heavy emphasis on advanced analytical skills and the ability to use programming languages such as R, Python and SQL for basic data mining. The ability to learn independently in a new domain and master complex code base is also listed as one of the key skills. A few of the postings list project management and organizational skills; however, this is not common.

When it comes to the knowledge required, the focus shifts from the technical domain to an understanding of fields such as sociology, critical data studies, and AI regulations. Many postings highlight that potential candidates need to be familiar with concepts such as AI/ML auditing, algorithmic impact assessments, assessment of fairness in predictive models, explainability, robustness, and human-AI interaction. Technical knowledge, such as understanding transfer-based language models and logistic regression model development, is also highlighted in the posting. Lastly, the job postings outline the need for strong interpersonal, verbal, and written communication skills. However, experience publishing and presenting at academic venues is not mentioned.

**Qualifications.** The majority of the job postings require a bachelor's degree in quantitative fields such as data science and computer science and prefer higher degrees (master's or Ph.D.). Companies are looking for candidates who have experience in data science, software engineering, and worked with large language models. Moreover, they are looking for experience in putting responsible AI principles into practice, evaluating the ethics of algorithms, and having basic familiarity with law and policy research. The ability and experience to translate AI ethic principles into practice are heavily emphasized throughout these job postings.

**Interpersonal Qualities.** The job postings emphasize the ability to work with people from different backgrounds. However, these job postings do not include a comprehensive list of values. A few postings mention being a self-starter, working collaboratively to resolve conflict, and caring deeply about the data used to train ML models as key attitudes. Being flexible, innovative, curious, adaptive, and passionate about tackling real word challenges are also some of the sought-after values.

## 4.3 Engineer

The engineer occupation is represented in 8 of the job postings. None of our interview participants belong to this category. The key responsibility of an engineer practicing AI ethics is to help establish a safety culture and system within an organization by developing technical tools. They are tasked with developing a workflow for modeling and testing for issues such as bias, explainability, safety, and alignment of AI systems. As part of this, engineers need to create code bases that could be used across the AI system development pipeline based on existing and evolving best practices.

**Skills and Qualifications.** Job postings for engineers place a significant emphasis on experience-based qualifications and skills. The companies represented in this dataset are looking for skills and experience in software development, dataset production pipelines, researching fairness and safety implications of ML systems, and the development of large language models. They are also looking for experience working in a fast-paced technology company. Based on these qualifications, the main set of skills are programming and AI/ML development skills and this needs to be supported by knowledge and familiarity with foundational concepts in AI/ML, fairness, explainability, system's safety, and safety life cycle management. Lastly, most of the job descriptions do not have a heavy emphasis on communication skills. Only a few mention excellent written and oral communication skills as a requirement.

*Interpersonal Qualities.* In contrast to the lack of emphasis on communication skills, these postings have a particular focus on the attitude and values of ideal candidates more so than any other occupation category. These attitudes include being result-oriented, willingness to contribute as needed (even if not specified) and keen to learn new concepts. They are looking for people who value working on challenging problems and care about the societal impact of their work.

#### 4.4 Researcher (Policy, design, science and technology studies, social sciences)

The second most frequent category of postings belongs to researchers that focus on topics such as policy, sociotechnical issues, and governance (14 job postings, 3 interviews). We created a separate group of positions as their responsibilities, skills, and qualities are sufficiently different from the technical researcher position. Candidates in this category need to conduct research, perform ethics or impact assessments of AI systems, act as a liaison and translator between research, product, policy, and legal teams, and lastly, advise on policy, standards, and regulations-related matters internally and externally. When conducting research, two different focus areas come up in the job postings: testing and evaluating AI system to inform policy and researching existing policies/regulations, and translating them into practice.

*Skills.* The job postings highlight two sets of distinct skills for this group of researchers. Firstly, these researchers require a basic level of programming, advanced analytics, and data visualization skills. Few positions highlighted the need for even more advanced ML and AI skills. It is noteworthy that despite these researchers' focus on policy, governance and sociotechnical issues, the postings still require them to have some data analytic skills. Secondly, these researchers need to have excellent facilitation, community-building, and stakeholder engagement skills. These two skills need to be complemented by strong leadership and management skills. The job postings heavily emphasize strong communication skills for this group of researchers. Besides the conventional skill of presenting and publishing papers, this group of researchers need to effectively work across different functionalities and disciplines. On a similar trend, these researchers need to have expertise in a variety of areas. They need to have a good understanding of "*qualitative and quantitative research methods*", reliably know the current and emerging "*legal and regulatory frameworks and policies*", be "*familiar with AI technology*" and have a good knowledge of practices, process, design, and development of AI technology. This is a vast range of expertise and often "*very difficult to recruit*" for as highlighted by our expert interviewees.

*Qualifications.* Just over half of the job postings list a Ph.D. in relevant areas as a requirement, including human-computer interaction, cognitive psychology, experimental psychology, digital anthropology, law, policy, and quantitative social sciences. Two postings require only a bachelor's or a master's in the listed areas. Similar to the technical researcher occupation, some positions do not specify any educational requirements and only focus on experience and skills. Our expert interviewees in this category are from a range of educational backgrounds ranging from a master's in

sociotechnical systems, a law degree combined with a background in statistics, and a master's in cognitive systems.

Besides experience in research, companies are looking for experience in translating research into design, technology development, and policy. A researcher explained that they need to do a lot of "translational work" between the academic conversation and product teams in companies. A good candidate for this occupation would have "*project management*", "*change management*", "*stakeholder engagement*", and "*applied ethics*" experience in a "fast-paced environment". All four of these skills do not appear in all of the job postings and interview discussions. However, a permutation of them appears throughout the job posting data and participants' responses.

*Interpersonal Qualities.* As emphasized strongly in both of the datasets, ideal candidates in this category need to have a "*figure-it-out somehow*" or "*make it happen*" attitude as explained by a participant. They are "*driven by curiosity and passion towards*" issues related to responsible AI development and are excited to engage with the product teams. Participants noted that ideal candidates in these roles are "*creative problem solvers*" who can work in a "*fast-changing environment*".

#### 4.5 Policy analyst

Policy analyst occupation is the least represented [1 expert interview, 4 job postings] in our data sources; however, considering the consistent list of competencies, we decided to include it within the proposed framework. The role of a policy analyst is to understand, analyze and implement a given policy within an organization. Moreover, they need to engage with policymakers and regulators and provide feedback on existing policies.

*Skills and Qualifications.* A policy analyst needs to have proven knowledge of laws, policies, regulations, and precedents applicable to a given technology when it comes to AI ethics-related issues. Moreover, all of the job postings highlight the importance of familiarity with AI technology. According to the job postings, a good candidate would have experience in interpreting policy and developing assessments for a given application. They also need to be skilled in management, team building, and mentorship. This finding echoes remarks from expert interviews. Even though none of the job postings specify an educational degree requirement, the expert we interviewed was a lawyer with a master's in technology law.

*Interpersonal Qualities.* The job postings in this category heavily emphasized values and attitudes. A good analyst needs to have sound judgment and outstanding personal integrity. They should be caring and knowledgeable about the impact of technology on society. Moreover, they enjoy working on complex multifaceted problems and are passionate about improving governance of AI systems. The expert interviewee's perspective closely matches these attributes. Participants elaborated that they needed to be "*brave*" and "*step up to ask questions and challenge status quo consistently over a long time*". As expected communications skills are considered critical for success. The expert interviewee significantly emphasized the importance of "*networking as a key factor*" in succeeding in their role.



## 4.6 Manager

We analyzed 7 management-related job postings and 5 expert interviewees in this category. The product managers take the role of incorporating responsible AI practices in the product development process. In contrast, program managers are often leading and launching a new program on establishing AI ethics practices within the organization. These programs often involve building an organization's capacity to manage responsible AI issues.

**Skills.** For both streams of management, the potential candidates need to have strong business acumen and a vision for the use/development of AI technology within an organization. Some of the key management skills highlighted in the job postings include the ability to manage multiple priorities and strategically remove potential blockers to success. Another sought-after skill is the ability to effectively engage stakeholders in the process. Expert interviewees also echoed the importance of this skill as their roles often involve getting people *"on board with new ways of thinking and creating"*. According to the job postings, good candidates for management need to have a practical understanding of the AI life cycle and be familiar with integrating responsible AI practices into a program or a product. Our interviewees note that they continuously need to *"learn and keep up with the fast-paced development of AI"*.

**Qualifications.** Not many postings have highlighted educational qualifications and instead focused on experience qualifications. However, the main educational qualification is a bachelor's degree with a preference for higher degrees. The postings have primarily highlighted a degree in a technical field such as computer science or software engineering. Interestingly the interviews reflect a different flavor of educational backgrounds. All of the experts we interviewed had at minimum a master's degree and the majority of them completed their studies in a non-technical field such as philosophy, media studies, and policy. However, these individuals had acquired a significant level of expertise in AI ethics through *"self-studying"* and *"engaging with the literature"* and the responsible AI *"community"*. For example, two of the participants trained in technical fields and had a significant level of industry experience. Similarly, they had learned about responsible AI through their own initiative.

On the other hand, the job postings heavily focus on *experience*, including a significant amount of technical know-how, experience focused on ML development, product and program management, and implementation of ethical and social responsibility practices within fast-paced technology companies. The interview participants had been *"working in the industry for some time"* before taking on these management roles. However, their range of experiences do not cover all of the required experiences outlined in the job descriptions. As expected, excellent communication skills are noted in the job descriptions and strongly echoed by the experts as well. The job postings do not necessarily elaborate on the nature of communication skills; however, the experts note that the *"ability to listen"*, *understand*, and sometimes *"persuade different stakeholders"* is key in such roles.

**Interpersonal Qualities.** Few of the job postings make remarks about attitudes/values and highlight that managers need to value

designing technology for social good and cooperation with other stakeholders. A good candidate for management should foster a growth mindset and approach their work with agility, creativity, and passion. All of the participants expressed their passion for developing ethical technology and indicate that they took a lot of initiative to learn and contribute to the field within their company and externally before they could take on their management roles.

## 4.7 Director

The job descriptions dataset has 4 postings for director positions and 2 of the expert interviewees have directorship roles. According to the job postings, director responsibilities include at least three of the following: a) lead the operationalization of AI ethics principles, b) provide strategic direction and roadmap towards enterprise-wide adoption and application of ethical principles and decision frameworks, and c) build internal capacity for AI ethics practice and governance. Depending on the nature of the organization and its need to incorporate AI ethics practices, these responsibilities vary in scope. For example, a director within a technology start-up will only be able to commit *"limited amount of time to operationalizing AI ethics principles and building internal capacity"* compared to a director within a larger technology company.

**Skills and qualifications.** According to the job postings, the key skill for being a director is having the ability to build a strong relationship with a broad community that helps define and promote best practice standards of AI ethics. An ideal director can effectively pair their technical skills/know-how with their management skills and policy/standards knowledge to develop strategic plans for the company. Experience in directing and leading teams, particularly in social responsibility practices within technology companies is highly valued for such positions. Only one job posting specifies an educational (a bachelor's related to policy development and implementation). Others only highlight experience. The two interviewees hold master's degrees in business and information systems respectively. They also had extensive industry experience that was not directly in AI ethics. However, their experience involved *"translation of policy within a technology application"*.

**Interpersonal Qualities.** As expected, according to the job postings a good candidate for directorship needs to have exceptional written and verbal communication skills, need to be able *"to articulate complex ideas"* to technical and non-technical audiences, *"engage and influence stakeholders"* and *"collaborate with people from different disciplines, and cultures"*. This set of skills was reflected in our expert interviews. Both interviewees emphasized how they maintain a good flow of communication with the employees and how they remain always open to having conversations on a needs basis. This allowed them to build trust within the company and pursue moving forward with their strategic plan. The job postings highlight the ability to earn trust in relationships as a sought-after value for a directorship role. A director should also be able to *challenge the status quo*, *be passionate about good technology development*, *be comfortable with ambiguity*, and *adapt rapidly to changing environment and demands*. Most importantly, a director needs to have *"a strong and clear commitment to the company values"* as they set the tone for others within the organization.



## 4.8 Emerging occupations

Besides the abovementioned classes of occupations, we found a few other positions that do not map easily to any of the existing categories. Considering the limited number of these positions, they do not justify a category of their own. However, we note these emerging roles to understand how they might shape up the responsible AI profession. These occupation titles include data ethicists (2 in job postings), AI ethics consultants (2 in interviews), dataset leads (2 in job postings), communication specialist (1 in job postings), safety specialist (1 in job posting) and UX designer (1 in job postings). The following describes the main function of these positions:

- Data ethicist: manage organizational efforts in operationalizing AI ethics practices through policy and technology development work. This role has similarities to the role of a policy analyst and data scientist.
- AI ethics consultant: apply their expertise in AI ethics to solve pain points for consulting clients.
- Dataset lead: curate datasets while accounting for fairness and bias-related issues.
- Safety specialist: use and test large language model-based systems to identify failures and errors.
- AI ethics communication specialist: write communication pieces that focus on AI ethics issues.
- UX designers: design user interfaces with ethics in mind.

## 4.9 Future of the responsible AI profession

Our interview participants shared a variety of responses to the question "what will the future of their job be like?". Some participants thought that eventually, *"everyone in a company will be responsible"* for understanding ethical and social issues of AI as part of their job. In this scenario, everyone would need to have the appropriate knowledge and skillset to apply responsible AI practices in their work or at least know when they need to ask for advice from internal or external experts.

On the contrary, many participants expressed that *"dedicated roles"* need to be recruited. These participants elaborate that recruitment for these roles is and will *"continue to be challenging"* as it is difficult to find people with interdisciplinary backgrounds and established industry work experience. Many of the managers we interviewed have chosen *"to build teams that come from different disciplinary backgrounds"* and provide *"professional development opportunities"* on the job. However, they also described that hiring people into these roles is challenging since corporate leaders are not always willing to invest a lot of resources in AI ethics. This often can lead to *"exhaustion and burn-out"* for individuals who currently hold these roles - this is especially true for small and medium-sized technology companies. According to participants, this will likely change with a progressive shift in the regulatory landscape.

## 5 DISCUSSION

Educators and employers play a pivotal role in shaping a responsible AI culture. In our efforts to create a competency framework that outlines the range of roles for responsible AI practitioners, we find that such frameworks can not only guide corporate leaders to recruit talent but also help grow their responsible AI capacity.

We find that the ability to work in an interdisciplinary environment, communicate and engage with diverse stakeholder groups, and the aptitude for curiosity and self-learning are consistently highlighted for all of the roles. This emphasizes the need to foster an environment where students and existing employees in different roles are encouraged to adopt interdisciplinary approaches/collaboration and explore responsible AI content.

In this section, we articulate how an interdisciplinary environment can be fostered, the importance of organizational support for responsible AI practitioners, and the need to proactively monitor the rapidly changing occupational demand and landscape for these roles.

### 5.1 Being able to work in an interdisciplinary environment is critical

Our results show that many of the responsible AI practitioners today come from non-traditional, non-linear, and interdisciplinary educational and work backgrounds to their current positions. The educational and work experiences of these participants span a multitude of fields and allowed them to develop a strong set of skills in navigating disciplinary boundaries and understanding problems from diverse perspectives. The participants often described their role as a *translator* and *facilitator* between different groups and disciplines within the organization. For instance, they remarked that a concept such as fairness, transparency, or ethically safe has completely different meanings depending on the personal and professional backgrounds of their audience. The participants often needed to translate what these concepts mean across different disciplinary boundaries (i.e. statistics and law).

Notably, while the job postings asked for a diverse array of skills and qualifications from multiple disciplines, those who hold such positions today are often specialized in one or two disciplines. However, they had been exposed to and worked across multiple disciplines in their professional career. The most important asset that our interviewees emphasized was being able to work across disciplinary boundaries. The candidates who successfully hold such positions are not *"ethical unicorn, full stack developers"* [45]. However, they have honed the skills necessary to translate and create solutions to responsible AI issues across multiple disciplines. Building on existing proposal to improve responsible AI practices [19, 35, 48] and education [26, 43], we posit that AI team leaders need to pay a special attention to hiring individuals with the capability to create, critique and communicate *across multiple disciplines*. Consider Furthermore, educators can get inspiration from education models in highly interdisciplinary fields such as healthcare and create curricula/spaces where students work with peers from different academic backgrounds [20, 29, 32].

### 5.2 Responsible AI practitioners are advocates - but they need organizational support

We find that responsible AI practitioners are often highly driven and motivated to make a positive impact. These individuals often hold a strong sense of valuing social justice and want to ensure that AI technology is developed in a way that is good for society's well-being. One of the most consistent ideas that came through in the interviews is the attitude that the participants had toward

their careers. Many of the interview participants took the time to immerse themselves in learning new topics and expressed that they were self-motivated to do so. This is especially true for the individuals who are taking some of these first positions in the industry. When looking at the career trajectory of many of the participants, we observe that they often created their own roles or came into a newly created role. Moreover, these individuals often needed to start their own projects and create relationships with others in the organization to measure their own progress and establish credibility.

Similar to any emerging profession many of the participants act as champions for ethical and safe development of AI. They are often working in an environment that questions and challenges the need for considering AI ethics principles. As some of the participants remarked, they often have to answer questions such as *"why do we need to pay for ethics assessments?"*, *"what is the value of considering AI ethics in a start-up?"*, or *"why should we put in the time? what is the value added?"*. This act of advocating for AI ethics is even more challenging when existing regulations do not have proper enforcement mechanisms for responsible AI practices [19]. Many of the participants assume the role of an advocate and often use their excellent communication skills to build relationships and capacity within their organization.

For the successful implementation of responsible AI practices, it is important that business leaders pay attention and support the advocacy efforts of these practitioners. Many of today's responsible AI practitioners are working with limited resources [40], have critical responsibilities [47], and are experiencing burn-out [30]. Whenever possible, leaders in AI companies need to create appropriate incentive structures, provide the necessary resources and communicate the value of establishing responsible AI practices to their employees so that these practitioners have the necessary support for the effective execution of their responsibilities. Recognizing the nature of these roles, educators can learn from existing methods [17, 22] and integrate leadership training into their curricula when addressing responsible AI-related content.

### 5.3 Educators and employers need to monitor and plan for the rapidly changing landscape of responsible AI roles

The nature of occupations in the AI industry is continually growing and shifting. The rapid technological development [3, 37], upcoming regulations [7] and global economic conditions [27, 33] impact how companies recruit and retain responsible AI expertise. Furthermore, there is a need for new educational efforts and programs for preparing new graduates to take on responsible AI practices. The proposed ontology provides a synthesis of roles that have emerged in responsible AI practice and it can serve as a planning tool for corporate leaders and educators.

Corporate leaders can use this ontology to build internal capacity for individuals who currently hold researcher, data scientist, engineer, policy advisor, manager, and director roles in their institutions. Depending on these companies' responsible AI needs and resources, business executives can work towards creating interdisciplinary teams for establishing responsible AI practice by recruiting individuals with the competencies outlined for each of

these roles. Besides recruiting and fostering for responsible AI competencies, these leaders need to communicate the importance of these practices and start by creating the appropriate organizational incentives and resources for adapting responsible AI practices. Government and non-governmental organizations could support such efforts, particularly small and medium-size companies, by formally recognizing such roles in their taxonomies of occupations [9, 21] and providing resources [12].

Current computer science and engineering education focuses primarily on teaching professional ethics [43]. There is minimal focus and resources on cultivating skills and knowledge required for cultivating the skills that focus on ethics in design [26]. On the other hand, there is a lack of clarity of how much students in social and political sciences need to work on their technical acumen to become skilled responsible AI practitioners [45]. Educators could use the list of competencies to develop a set of learning objectives and examine the efficacy of different teaching pedagogies in supporting these objectives. Moreover, Educators can use the competency framework as a tool for acquiring resources for further curricula and program development.

Notably, the proposed ontology primarily focuses on type of roles, responsibilities and skills without addressing other important factors in recruitment and education efforts such as diversity of individual who get to learn about responsible AI issues or take such roles in the industry. Therefore, it is critical that users of this ontology, consider factors that are not captured in the scope of this ontology. Furthermore, considering the rapidly changing conversation around responsible AI practices, the type of roles in this ontology will shift and expand. We invite the community of researchers, practitioners and educators to reflect on these roles and build on this ontology.

## 6 CONCLUSION

With the increased regulatory activities in the industry, companies have the incentive to ensure responsible AI development. In this work, we found seven different type of roles and their corresponding responsibilities, skills, qualifications, and interpersonal qualities expected in today's responsible AI practitioner. We propose a preliminary competency framework for responsible AI practitioners and highlight the importance of creating interdisciplinary teams and providing adequate organizational support for individuals in these roles.

## ACKNOWLEDGMENTS

We thank our study participants for taking the time to share their experiences, expertise, and feedback. We also thank our anonymous reviewers for their deep engagement and valuable feedback on this paper. This work benefitted greatly from the data collection and analysis assistance from our collaborators Sandi Mak, Ivan Ivanov, Aidan Doudeau, and Nandita Jayd at Vanier College, Montreal. We are grateful for their contributions. Finally, this work was financially supported by the Natural Sciences and Engineering Research Council of Canada and Pôle montréalais d'enseignement supérieur en intelligence artificielle.

## REFERENCES

- [1] 2021. A European approach to artificial intelligence - shaping Europe's digital future. (2021). <https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>
- [2] 2021. Competence and competency frameworks - factsheets - CIPD. <https://www.cipd.co.uk/knowledge/fundamentals/people/performance/competency-factsheet>
- [3] 2021. Embracing the rapid pace of AI. *MIT Technology Review* (May 2021).
- [4] 2022. AI Ethics Toolkit. (2022). <https://www.ibm.com/artificial-intelligence/ethics>
- [5] 2022. Algorithmic Impact Assessment. <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html>
- [6] 2022. IEEE Ethics in Action in Autonomous and Intelligent Systems.
- [7] 2022. The Regulation of Artificial Intelligence in Canada and Abroad: Comparing the Proposed AIDA and EU AI Act. <https://www.fasken.com/en/knowledge/2022/10/18-the-regulation-of-artificial-intelligence-in-canada-and-abroad>. Accessed: 2023-2-15.
- [8] 2022. Responsible AI resources. (2022). <https://www.microsoft.com/en-us/ai/responsible-ai-resources>
- [9] Employment & Training Administration. 2022. O\*NET database Content Model. <https://www.onetcenter.org/content.html>
- [10] Ethical AI Advisory. 2023. Ethical AI Advisory. <https://www.ethicalai.ai/>. Accessed: 2023-3-15.
- [11] David Anderson, Joy Bonaguro, Miriam McKinney, Andrew Nicklin, and Jane Wiseman. 2020. Ethics and algorithms toolkit. <https://ethicstoolkit.ai/>
- [12] James Bessen, Stephen Michael Impink, and Robert Seamans. 2022. The Cost of Ethical AI Development for AI Startups. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (Oxford, United Kingdom) (AIES '22). Association for Computing Machinery, New York, NY, USA, 92–106.
- [13] Sherry Blok, Joel Trudeau, and Robert Cassidy. 2021. *Artificial Intelligence Competency Framework Table of Contents*. Technical Report September.
- [14] Jason Borenstein and Ayanna Howard. 2020. Emerging challenges in AI and the need for AI ethics education. *AI Ethics* 1, 1 (2020), 61–65. <https://doi.org/10.1007/s43681-020-00002-7>
- [15] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* 3, 2 (Jan. 2006), 77–101.
- [16] Frédéric Bruneault, Cégep André-laurendeau, Andréane Sabourin Laflamme, Gabrielle Fillion, Neila Abtroun, and Andrew Freeman. 2022. *AI Ethics Training in Higher Education: A Competency Framework*. Technical Report February.
- [17] Lori Carter. 2011. Ideas for adding soft skills education to service learning and capstone courses for computer science students. In *Proceedings of the 42nd ACM technical symposium on Computer science education* (Dallas, TX, USA) (SIGCSE '11). Association for Computing Machinery, New York, NY, USA, 517–522.
- [18] ORCAA Consulting. 2023. ORCAA. <https://orcaarisk.com/>. Accessed: 2023-3-15.
- [19] Sasha Costanza-Chock, Inioluwa Deborah Raji, and Joy Buolamwini. 2022. Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 1571–1583.
- [20] Jean A Dyer. 2003. Multidisciplinary, Interdisciplinary, and Transdisciplinary Educational Models and Nursing Education. *Nurs. Educ. Perspect.* 24, 4 (2003), 186.
- [21] European Commission. 2022. The ESCO Classification. <https://esco.ec.europa.eu/en/classification>
- [22] John V Farr and Donna M Brazil. 2009. Leadership Skills Development for Engineers. *Engineering Management Journal* 21, 1 (March 2009), 3–8.
- [23] Jessica Fjeld, Nele Achten, Hannah Hillgoss, Adam Nagy, and Madhulika Sri Kumar. 2020. Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI. (Jan. 2020).
- [24] Heidi Furey and Fred Martin. 2019. AI education matters: a modular approach to AI ethics education. *AI Matters* 4, 4 (Jan. 2019), 13–15.
- [25] Olivia Gambelin. 2020. Brave: what it means to be an AI Ethicist. *AI Ethics* 1, 1 (2020), 87–91. <https://doi.org/10.1007/s43681-020-00020-5>
- [26] Natalie Garrett, Nathan Beard, and Casey Fiesler. 2020. More than “if time allows”: The Role of Ethics in AI Education. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (New York NY USA). ACM, New York, NY, USA.
- [27] Sharon Goldman. 2022. Why Meta and Twitter's AI and ML layoffs matter. <https://venturebeat.com/ai/why-meta-and-twitters-ai-and-ml-layoffs-matter-the-ai-beat/>. Accessed: 2023-3-12.
- [28] Radhika Gorur, Leonard Hoon, and Emma Kowal. 2020. Computer Science Ethics Education in Australia—A Work in Progress. In *2020 IEEE International Conference on Teaching, Assessment, and Learning for Engineering (TALE)*. IEEE Explore, 945–947.
- [29] P Hall and L Weaver. 2001. Interdisciplinary education and teamwork: a long and winding road. *Med. Educ.* 35, 9 (Sept. 2001), 867–875.
- [30] Melissa Heikkilä. 2022. Responsible AI has a burnout problem. *MIT Technology Review* (Oct. 2022).
- [31] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* 1, 9 (2019), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- [32] Renate G Klaassen. 2018. Interdisciplinary education: a case study. *Eur. J. Eng. Educ.* 43, 6 (Nov. 2018), 842–859.
- [33] Will Knight. 2022. Elon Musk Has Fired Twitter's 'Ethical AI' Team. *Wired* (Nov. 2022).
- [34] AI Ethics Lab. 2019. AI Ethics Lab. <https://aiethicslab.com/>. Accessed: 2023-3-15.
- [35] Michael A. Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Comput. Hum. Interact.* 1–14. <https://doi.org/10.1145/3313831.3376445>
- [36] Matti Mäntymäki, Matti Minkinen, Teemu Birkstedt, and Mika Viljanen. 2022. Putting AI Ethics into Practice: The Hourglass Model of Organizational AI Governance. (June 2022). arXiv:2206.00335 [cs.AI]
- [37] Taylor Meek, Husam Barham, Nader Beltaif, Amani Kaadoor, and Tanzila Akhter. 2016. Managing the ethical and risk implications of rapid advances in artificial intelligence: A literature review. In *2016 Portland International Conference on Management of Engineering and Technology (PICMET)*. IEEE Explore, 682–693.
- [38] Jakob Mökander, Maria Axente, Federico Casolari, and Luciano Floridi. 2022. Conformity Assessments and Post-market Monitoring: A Guide to the Role of Auditing in the Proposed European AI Regulation. *Minds Mach.* 32, 2 (2022), 241–268.
- [39] Ajung Moon, Shalaleh Rismani, Jason Millar, Terralynn Forsyth, Jordan Eshpeter, Muhammad Jaffar, and Anh Phan. 2019. Foresight into AI Ethics.
- [40] Emanuel Moss and Jacob Metcalf. 2020. *Ethics Owners: A new model of organizational responsibility in data-driven technology companies*. Technical Report. 1–71 pages.
- [41] Ehsan Nabavi and Chris Browne. 2023. Leverage zones in Responsible AI: towards a systems thinking conceptualization. *Humanities and Social Sciences Communications* 10, 1 (March 2023), 1–9.
- [42] Judith Simon Pak-Hang Wong. 2020. *Thinking About 'Ethics' in the Ethics of AI*. Technical Report.
- [43] Tina L Peterson, Rodrigo Ferreira, and Moshe Y Vardi. 2023. Abstracted Power and Responsibility in Computer Science Ethics Education. *IEEE Transactions on Technology and Society* (2023), 1–1.
- [44] Thomas P Quinn and Simon Coghlan. 2021. Readying Medical Students for Medical AI: The Need to Embed AI Ethics Education. (Sept. 2021). arXiv:2109.02866 [cs.AI]
- [45] Inioluwa Deborah Raji, Morgan Klaus Scheuerman, and Razvan Amironesei. 2021. “you can't sit with us”: Exclusionary pedagogy in AI ethics education. In *FAccT 2021 - Proc. 2021 ACM Conf. Fairness, Accountability, Transpar.* 515–525. <https://doi.org/10.1145/3442188.3445914>
- [46] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *FAT\* 2020 - Proc. 2020 Conf. Fairness, Accountability, Transpar.* (2020), 33–44. <https://doi.org/10.1145/3351095.3372873> arXiv:2001.00973
- [47] Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2021. Where Responsible AI meets Reality: Practitioner Perspectives on Enablers for Shifting Organizational Practices. *Proc. ACM Human-Computer Interact.* 5, CSCW1 (2021), 1–23. <https://doi.org/10.1145/3449081> arXiv:2006.12358
- [48] Shalaleh Rismani, Renee Shelby, Andrew Smart, Edgar Jatho, Josh A Kroll, Ajung Moon, and Negar Rostamzadeh. 2023. From Plane Crashes to Algorithmic Harm: Applicability of Safety Engineering Frameworks for Responsible ML. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (CHI '23) (Hamburg, Germany), Vol. 1. Association for Computing Machinery.
- [49] Seema Sanghi. 2016. *The handbook of competency mapping* (3e ed.). SAGE.
- [50] Daniel Schiff, Bogdana Rakova, Aladdin Ayeshe, Anat Fanti, and Michael Lennon. 2021. Explaining the Principles to Practices Gap in AI. *IEEE Technol. Soc. Mag.* 40, 2 (June 2021), 81–94.
- [51] Renee Shelby, Shalaleh Rismani, Kathryn Henne, Ajung Moon, Negar Rostamzadeh, Paul Nicholas, N'mah Yilla, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. 2022. Identifying Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. (Oct. 2022). arXiv:2210.05791 [cs.HC]
- [52] Mona Sloane and Janina Zakrzewski. 2022. German AI Start-Ups and “AI Ethics”: Using A Social Practice Lens for Assessing and Implementing Socio-Technical Innovation. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 935–947.
- [53] L.M. Spencer and S.M. Spencer. 1993. *Competence at Work - Models for Superior Performance*. John Wiley & Sons, Inc., New York.
- [54] Kees Stuurman and Eric Lachaud. 2022. Regulating AI. A label to complete the proposed Act on Artificial Intelligence. *Computer Law & Security Review* 44 (April

- 2022), 105657.
- [55] Qiaosi Wang, Michael Adam Madaio, Shivani Kapania, Shaun Kane, Michael Terry, and Lauren Wilcox. 2023. Designing responsible AI: Adaptations of UX practice to meet responsible AI challenges. (2023).
  - [56] Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. Taxonomy of Risks posed by Language Models. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 214–229.
  - [57] David Gray Widder and Dawn Nafus. 2023. Dislocated accountabilities in the “AI supply chain”: Modularity and developers’ notions of responsibility. *Big Data & Society* 10, 1 (Jan. 2023), 20539517231177620.
  - [58] Randi Williams. 2021. How to Train Your Robot: Project-Based AI and Ethics Education for Middle School Classrooms. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education* (Virtual Event, USA) (SIGCSE '21). Association for Computing Machinery, New York, NY, USA, 1382.