How Can Socio-Technical Systems Design Approaches Ensure Autonomy in Hiring Practices?

Liam Kedem

April 2021

Abstract

Artificial Intelligence (AI) is widely used in hiring practices to identify the most suitable candidate for a vacancy. This is due to the promise of higher overall efficiency and lower costs. However, these AI-powered tools may create an inaccurate conception of the applicant's suitability to the vacancy by numerically quantifying context-dependent variables. If only this inaccurate conception is used to judge an applicant, this is a violation of the applicant's autonomy over their self-representation. This paper argues that such a problem could be solved by adopting a broader design scope - *Socio-Technical Systems Design* (STSD). STSD approaches have not been widely applied to AI or hiring practices yet. Therefore the main contribution of this paper is to bridge this gap by exploring possible STSD approaches which can be applied to ensure the applicants' autonomy over self-representation. This paper suggests combining methodologies and design principles from two STSD approaches - *Design for Values* and *Systemic Design*. The findings from Design for Values suggest that key stakeholders should be involved in the design process. Therefore, the designers should conduct a stakeholder analysis to identify the key stakeholders, followed by an investigation to explore the stakeholders' needs, and the values which the proposed system could implicate. Systemic Design offers design principles that should be utilized during the design process. These principles consist of: expanding the problem space; focus on the relationship between the key stakeholders; and follow an iterative, experimental, and evolutionary design approach throughout the design process. This nuanced, stakeholder-centric approach results in an inclusive, transparent, multi-discipline, and socially aware process which is necessary to understand the complex social context, and its socio-ethical issues.

1 Introduction

In recent years artificial intelligence (AI) has started to be integrated into many different social domains, including criminal justice [Yong, 2018] and health care [Daley, 2019], due to the promise of data-driven and automatic decisions leading to higher overall efficiency. Given that AI is such a broad and vague term, this paper focuses specifically on machine learning tools, which AI Myths [n.d.] define as "an approach to AI that relies on training algorithms on large datasets so that they develop their own rules" (p. 1). This form of AI is widely used in hiring practices [Albert, 2019], which is the focus of this paper.

The hiring process consists of two phases, the recruitment phase, and the selection phase. Starting with the recruitment phase, an organization wants to attract quality candidates to fill an existing vacancy. Historically, this was done through newspaper job posting, but the introduction of the internet has provided new possibilities for advertising job openings, such as social media (e.g. LinkedIn), professional associations [Levin, 2016] and career-related websites (e.g. Magnet.me). Through this process, hiring managers can reach out to potential job applicants and encourage them to apply for available vacancies.

This is followed by the selection phase, which is the focus of this paper. During this phase the talent acquisition specialist is responsible for pre-screening, reviewing, prioritizing, and shortlisting applicants. The goal of this process is to identify the most suitable candidate for the job by assessing the candidate's behavioral traits, attitude, and domain knowledge to choose the most suitable candidate to fill the vacancy. This has been traditionally done by an in-person interview or test [Singh, 2020].

Albert [2019] identified 11 AI-powered tools used across the recruitment and selection phases. These AI tools are mostly being adopted by larger, tech-focused, and/or innovative firms, such as IBM, LinkedIn, and PwC.

After conducting research, by for instance visiting the websites of the vendors of these AI tools, I have identified three tools that can potentially use numerical measurements to assess an applicant. These tools are:

- CV screening software which assesses each candidate according to data from their CV and ranks them according to their fit to the job and organization. This creates a list of all the applicants, ranking them from best to worst fit.
- AI-powered psychometric tests, which are tests trying to assess who the candidate is and what their knowledge level and/or personality is.
- Video screening software, which analyses video interviews to assess the applicant's fit to the job and organization.

These are three completely different assessment tools with the common goal of picking out the most suitable candidate for an open vacancy.

Additionally, these tools are all meant to solve certain problems in hiring practices. For instance, the goal of CV screening software is to instantly review a large amount of CVs, filter out and rank them according to certain criteria. The alternative is for humans to do this manually, but this process is time-consuming, expensive, and the rate of errors increases as the amount of CVs increases. The various vendors of this type of tool claim that using this tool will reduce bias, cost, and/or issues due to human fatigue while improving diversity and allowing recruiters to focus on more essential tasks [Albert, 2019].

Although these benefits sound very promising, it is important to keep the various socio-ethical challenges raised by the use of AI in mind, such as discrimination [Shin, 2020], invasion of privacy [Kerry, 2020] and social inequality [Caestano and Simpson-Young, 2021].

The focus of this paper is the individual's autonomy over self-representation. Autonomy accumulated different meanings over time and it is context-dependent, but Sensen [2012] defined autonomy in the eyes of Immanuel Kant as "The moral right one possesses, or the capacity we have in order to think and make decisions for oneself providing some degree of control or power over the events that unfold within one's everyday life". This means autonomy is an individual's capacity for self-governance or self-determination [Dryden]. The reason autonomy over self-representation is so important, is because it can be seen as a key dimension of human dignity [Aizenberg and van den Hoven, 2020, Halbertal, 2015]. Furthermore, human dignity is one of the core values of human rights [Official Journal of the European Union, 2012]. Therefore violating one's autonomy over self-representation can be seen as a violation of one's fundamental rights, their dignity.

An example where an individual's autonomy over self-representation is violated can be exemplified by the situation where a CV screening software, as described above, ranks a vacancy's applicants. The software assesses every applicant's CV and gives them a score based on the applicant's fit for a certain job and/or organization. In general, machine learning algorithms do this by assigning numerical values to a set amount of target variables. However, the target variables in hiring practices are context-dependent, making them highly fluid and contestable. Therefore, when the algorithm tries to quantify these complex variables, it ends up removing the context and contestability from these variables [Aizenberg and van den Hoven, 2020, Delandshere and Petrosky, 1998]. This results in an inaccurate conception of the applicant. The applicant's autonomy over self-representation is being violated when the algorithm's inaccurate conception of the applicant becomes the new focus of the hiring organisation. This leads to the applicant being judged based on the algorithm's conception of the applicant rather than the applicant's true competencies. This situation is called data determinism [Broeders et al., 2017].

This paper tackles this issue by broadening the design scope from a techno-centric approach to a socio-technical approach. Socio-technical systems acknowledge the importance of human, social, organizational, and technical factors [Baxter and Sommerville, 2011]. In contrast, techno-centric systems mainly consider the technical factors.

Socio-technical systems design (STSD) approaches have not been widely applied in AI and hiring practices yet. The goal and contribution of this paper are to bridge this gap by exploring how various STSD methods can be applied to ensure individuals' autonomy over self-representation.

The main research question is: *How can socio-technical systems design approaches ensure autonomy in hiring practices?*. The methodology for answering this question is covered in section 2, followed by delving deeper into data determinism (section 3) and its root cause in machine learning applications (section 4). Next, how and why to apply STSD approaches are covered; followed by an exploration of existing STSD methodologies that can be applied in the design process to address data determinism (section 5). A reflection on the reproducibility of this paper is provided in section 6. Finally, section 7 summarises the main findings in this paper, analyses the multidisciplinary nature of the design team, and discusses the limitations and direction for future work.

2 Methodology

A literature study is conducted in order to answer the main research question. The goal of the literature study is to get an in-depth understanding of the problem and analyze possible solutions for it. That means finding literature on data determinism, its possible causes, and different STSD approaches which can help solve the problems. The literature is mainly gathered through Google for background information and Google Scholar for scientific and review papers. More specifically, the references in this paper are either: suggested by my supervisor Evgeni Aizenberg, gathered by combining relevant keywords—such as data discrimination, artificial intelligence, socio-technical systems, etc.—into a query, or by following relevant references from other papers.

3 Data Determinism

AI generally works by combining a large amount of data and learning from patterns or features hidden in the data [Sas, 2021]. It is important to note that predictions made by these algorithms are usually based on statistical correlations rather than causal evidence [Kitchin, 2014]. This means that a person's score is dependent on the input data of the algorithm and the correlations the algorithm finds between the applicant and the rest of the data. This is a case of data determinism [Broeders et al., 2017], which is

"a situation in which data collected about individuals and the statistical inferences drawn from that data are used to judge people on the basis of what the algorithm says they might do/be/become, rather than what the person actually has done and intends to do." [Aizenberg and van den Hoven, 2020, p .7]

This directly compromises the individual's autonomy over self-representation, since the individual is not in control anymore of how they are being represented, it is based on the algorithm's judgment. This is especially a problem if the algorithm is non-transparent and/or uncontestable.

In the former case, this could lead to a state of helplessness since the applicant is not able to know why they are getting a certain score, violating the individual's human dignity [Halbertal, 2015]. This does not help the organization either, since they do not get clarification on why an applicant gets a certain score, and therefore what makes an applicant viable.

In the latter case, if an applicant is aware that an algorithm is assessing their fit they might adapt their behavior/views/achievements in a way that would comply with the algorithm's view of a good candidate. This means that instead of being one's true self, the applicant's personality is compromised to achieve a desired outcome. This is disadvantageous for both parties involved, since the applicant can not act as them-self and the organization might hire a person which does not match their CV.

Data determinism is closely related to information self-determination, which is defined as the right of an individual to decide what information about themselves should be communicated to others and under what circumstances [Westin, 1970]. This might not be implicated directly by the CV screening software mentioned above since the individual is actively sharing their CV with a certain organization and consenting for its data to be used in this process. But this is not the case with all profiling algorithms. The increase in the pervasiveness of profiling algorithms, which base their outputs on big data sets, have major implications on the individuals' privacy and transparency, and therefore the individual's autonomy over self-representation [Aizenberg and van den Hoven, 2020]. This is because privacy is essential for an individual to behave as their true self while being protected from having their personality distorted by a profiling algorithm dissecting their data for information such as hometown, education, and likes on Facebook [Matz et al., 2017].

This section defined and delved deeper into data determinism. The next section analyses a major underlying factor causing data determinism, especially in machine learning applications. This is the use of numerical measurements and their limitations in quantifying complex and context-dependent behavior.

4 Limitations of Numerical Measurements

Using numerical measurements has various benefits. First of all, numerical measurements are useful for representing simple and discrete behaviors that occur consistently across different individuals, contexts, and times. Secondly, they are very efficient forms of measurement. It is certainly easier and less time-consuming to score someone on a point system rather than explicitly have to articulate the evidence each rating is based on. Thirdly, after assigning the scores they can easily and economically be recorded and administered. Finally, and arguably most importantly, the public trusts the quality and fairness of quantified judgments, leading to implicitly supporting this measurement [Delandshere and Petrosky, 1998]. This leads to the general assumption that numerical measurements are more easily interpretable, defendable and comparable than other assessment forms; a score can be higher or lower and can easily be compared to other participant's scores or a set passing score [Delandshere and Petrosky, 1998].

Given these advantages of using numerical measurements, it should still be questioned whether they ought to be utilized to quantify complex events.

Generally, a machine-learning algorithm assigns a numerical value to a set amount of target variables, by using the given input data. These values are then used to essentially create a conception of the applicant [Aizenberg and van den Hoven, 2020]. Problems arise when the target variables are highly fluid, contextual, and contestable, which is indeed the case in hiring practices (e.g. An AI-powered psychometric test trying to quantify an individual's personality trait). Aizenberg and van den Hoven [2020] argue that algorithms try to create a conception of these target variables by reducing them to a set of quantifiable inputs and outputs, thereby removing the fluidity, context, and contestability from these variables. This results in the applicant's autonomy over self-representation being at risk since the algorithms' inaccurate conception of the applicant becomes the new focus of the hiring organization, rather than the applicant's actual competencies, leading to data determinism.

Therefore it is important to consider the intended purpose of these assessments and critically question whether it is even reasonable to numerically quantify these complex variables, such as personality traits, in the first place. Delandshere and Petrosky [1998] argue that these broad and complex variables lose meaning when they are forced to be sorted into numerical categories without actually being interpreted, especially since they are being taken out of context. These broad and complex personality traits, such as creativity, show in different ways according to the specific situation at hand. These complex performances can not be assessed by one specific instance of creativity, they should be measured across different situations over time. This means one can not infer the total domain of someone's personality trait from one numerical value. This is too narrow, restrictive and should not be seen as a representative consistent sample of one's universal personality.

Because of these given reasons, designers should explore the different ways to assess and measure the given variables rather than implicitly supporting numerical measurements without questioning their validity.

This section covered the use of numerical measurements in machine learning applications, and how this leads to data determinism when applied in hiring practices. The following section expands the design scope from a techno-centric approach to a socio-technical approach, to combat data determinism.

5 Socio-Technical Systems Design

Data determinism—in AI-based hiring tools—is mainly caused by following a techno-centric design approach. Selbst et al. [2019] describe this as the *solutionism trap*, which occurs when one does not consider the possibility that the best solution for the problem at hand may not involve technology at all. This means a system might meet all the set technical requirements, but it could still be a failure due to not properly considering additional factors, such as the humans affected the system [Baxter and Sommerville, 2011]. Selbst et al. [2019] describe this as the *framing trap*, which arises by failing to model the entire system correctly, due to putting too much focus on the algorithm's performance and not questioning the validity of one's (often implicit) abstraction choices.

A broader design scope is therefore required, which socio-technical systems design (STSD) offers. These systems acknowledge the importance of the complex interaction between humans, machines, and environmental aspects of the system, to make any technology work as intended [Emery and Trist, 1960]. Therefore STSD considers human, social, organizational, and technical factors when designing socio-technical systems [Baxter and Sommerville, 2011]. This means both the technical and social subsystems are considered together, as one integrated whole [Bednar and Welch, 2020].

Failure to focus on both of these systems simultaneously, by exclusively focusing on one while excluding the other, will likely degrade the overall performance and utility of the system [Baxter and Sommerville, 2011]. However, applying these design methods effectively result in a better understanding of the interaction between the human, social and organizational factors. This interaction dictates the work practices in the organization and the technical systems used for assistance. Additionally, it is shown that opting for STSD delivers more value to stakeholders while creating more acceptable systems to the end-users [Baxter and Sommerville, 2011].

The goals of any socio-technical system can be achieved by various means, therefore STSD methods will generally not provide a system designer with a detailed process they should follow [Baxter and Sommerville, 2011]. van der Bijl-Brouwer and Malcolm [2020] explain that this is the case because of the high level of complexity and unpredictability of socio-technical systems; pre-determined solutions simply aren't sufficient to cover the whole system. This goes against the deeply ingrained notion most computer scientists get in their training since task-centered abstraction allows the same solution to be applied across multiple settings; Code transferability is even glorified as a core value within the fair machine learning community [Selbst et al., 2019]. Selbst et al. [2019] warn that this could lead to the *portability trap*, which occurs when one might fail to realize that an algorithm in one social context might be inaccurate, or even do harm when applied to a different context.

As mentioned in section 4, abstraction is necessary for machine learning algorithms. However, abstracting too much from context-dependent variables ends up stripping away the social context, leading to it not being accounted for as it should be. This is exactly what we want to prevent by using STSD approaches. Selbst et al. [2019] warn the designers of these systems to be aware of five possible abstraction traps. The solutionism, framing, and portability traps have been covered above. Therefore two traps remain.

The *formalism trap* concerns the definition of the underlying value one tries to ensure with their system (such as autonomy or fairness). These social concepts cannot be fully accounted for through mathematical formulations, since they can be highly procedural, context-dependent, and/or contestable. Simplifying these definitions into a mathematical formulation can not possibly capture the full range of meanings these complex values can behold.

Next, the *ripple effect* concerns the shift in societal values present in the pre-existing social system due to the introduction of the new technology. It is important to understand how the technology interacts with the social system already in place and monitor whether there is a possibility this will alter the social values in the system.

Selbst et al. [2019] argue that one should fully spell out these five traps when considering a technical solution. That means going over the traps and determining whether a technical solution is appropriate for the social context at hand. This requires a nuanced understanding of the relevant social context, its politics, and its stakeholders.

The next subsections describe two STSD approaches, Design for Values (5.1) and Systemic Design (5.2). These approaches have been chosen due to their stakeholder/value-centric views which are necessary to get an inclusive and socially aware design process.

5.1 Design for Values

Design for Values refers to the explicit translation of vague social values into concrete context-dependent design requirements [van den Hoven et al., 2015]. It acts as an umbrella term incorporating various design methodologies, such as Value Sensitive Design [Friedman, 1997] and Participatory Design [Schuler and Namioka, 1993].

Translating abstract values into concrete design requirements is a core process within Design for Values, referred to as *value specification* [van de Poel, 2013]. First, the abstract values are expanded into norms which are "properties or capabilities that the designed technology should exhibit in order to support desired values " [Aizenberg and van den Hoven, 2020, p.3]. These norms are then further specified as concrete socio-technical design requirements. An example of this is taking autonomy over self-representation as a core value. One of its' norms might be the "ability to change". This norm can be further specified as a concrete design requirement which states that the user should be able to change their personal information [ter Haar Romenij, 2020].

This process results in a structured and transparent model called a value hierarchy, where one can see the design requirements that the technology has to adhere to in order to support the core social values. An important aspect of this process is that it is transparent, making it up for discussion. The key stakeholders may disagree with the proposed decomposition of the values and norms, which can then be followed by a focused debate. This nuanced, context-dependent, and stakeholder-centric approach results in an inclusive and socially aware conversation which is necessary in order to understand the complex social context and its socio-ethical issues.

Next, two methodologies incorporated in Design for Values are covered, namely *Value Sensitive Design* and *Participatory Design*.

5.1.1 Value Sensitive Design

Value Sensitive Design [Friedman et al., 2006] is seen as the historical origin of centralizing social values in the design process of new technology [van den Hoven et al., 2015]. Friedman et al. [2006] propose conducting three types of investigations in the design process.

> *Conceptual investigation* identifies direct stakeholders, indirect stakeholders, and the values implicated by the use of technology.
>
> *Empirical investigation* explores the stakeholders' needs, understandings, and experiences in relation to the technology and implicated values.
>
> *Technical investigation* concerns the technology's specific features. This includes designing new technology to support particular values and analyzing how existing technologies implicate values in a social context.

These three investigations are iterative and integrative. Therefore they are meant to inform each other. One should not view these as separate investigations since they could overlap or be performed in different orders [van den Hoven et al., 2015].

Friedman et al. [2006] laid out 14 unique Value Sensitive Design methods, which can be used in different scenarios. The following part of this section highlights some of these methods, starting with *Stakeholder Analysis* due to the importance of being able to identify the key stakeholders in a social system. Next, two stakeholder engagement techniques—*Envisioning Cards* and *Value Dams and Flows*—are highlighted to get a deeper understanding of the stakeholders' needs, social context, and implicated values.

Stakeholder Analysis Clegg [2000] claims that the design of a system should reflect the needs of the stakeholders. A stakeholder is defined by Fernando [2021] as "a party that has an interest in a company and can either affect or be affected by the business" (p. 1). These are then further divided into direct and indirect stakeholders. Friedman et al. [2006] define direct stakeholders as individuals or organizations who directly interact with a system or its output, and indirect stakeholders as individuals/organizations who are affected by the use of the system. Design for Values makes this distinction since it is important to consider everyone who is significantly affected by the technology. This means one has to look further than only the clients and users, for example [Friedman et al., 2006].

The goal of a stakeholder analysis is to identify the direct and indirect stakeholders in a social context and identify for each stakeholder what the potential benefits and harms are due to the proposed system [van den Hoven et al., 2015].

Generally, there are two key stakeholder groups involved in the selection process. First, there are the applicants, and secondly, there is the hiring organisation.

To be more specific, a stakeholder analysis has to be conducted to discover and select the key stakeholders specific to one's hiring organisation and context. A large organisation might have a dedicated hiring manager, team leader, and recruitment team which could all be key stakeholders [Steven, 2020]. On the other hand, a startup might only have a handful of employees, possibly leading to everyone being responsible for selecting a candidate. It is due to these variations that a stakeholder analysis should be conducted for one's context.

An important critique towards Value Sensitive Design is that it does not offer any systematic or comprehensive stakeholder analysis technique [University of Siegen and Yetim, 2011]. However, stakeholder analysis is not unique to Value Sensitive Design [van den Hoven et al., 2015]. This method is also used in other fields such as public policy, conflict resolution, and business administration. Therefore techniques from those fields can be used, such as the influence-impact grid, salience model, and power versus interest grid [Gudavajhala, 2017]. All of these techniques are designed to map out the various stakeholders in a system and understand the level of interest, power, and/or influence each stakeholder has within the social context.

It should be noted that not all other areas which use stakeholder analysis have the concept of indirect stakeholders. This is an important aspect of using Value Sensitive Design, which should be reflected in the used stakeholder analysis technique.

Finally, Christakis and Bausch [2006] argue that a principle called requisite variety should be applied in order to ensure an optimal selection of stakeholders in this multi-stakeholder design process. Requisite variety boils down to making sure the variety within a social system is well represented by the stakeholders. This is needed for the dialogue to lead to valid and effective resolutions. One should keep this in mind when conducting a stakeholder analysis.

Stakeholder Engagement After identifying and assembling the key stakeholders an empirical investigation should be conducted to explore the stakeholders' needs, understandings, concerns, and experiences with respect to the technology and the values it might implicate. There are a wide range of methods available in order to elicit these requirements, spanning various design methodologies such as Participatory Design (section 5.1.2) and Requirements

Engineering. Zowghi and Coulin [2005] identified eight Requirements Engineering techniques used for requirement elicitation, namely: interviews, domain analysis, group-work, ethnography, prototyping, goal-based approaches, scenarios, and viewpoints. It should be noted that although these techniques involve stakeholders and their needs, only a few explicitly focus on both direct and indirect stakeholders and their values [Detweiler and Harbers, 2014]. Therefore these methods from Requirements Engineering could be used with a Participatory Design point of view (section 5.1.2), requiring both designers and engineers to collaborate. Value Sensitive Design also offers various stakeholder engagement techniques, such as *envisioning cards* and *value dams and flows*.

Envisioning Cards are used to evoke consideration and discussion surrounding topics such as values within an interactive system and the long-term/indirect effects of proposed solution(s). This set of cards is based on the four "envisioning criteria": stakeholders, time, values, and pervasiveness [Nathan et al., 2008]. Each card highlights one of the four envisioning criteria, designed to provoke values-oriented reflection throughout the design, implementation, and evaluation of a given system [Friedman and Hendry, 2012]. By understanding each stakeholders' values and how they affect the values of other stakeholders, the designers get a deeper understanding of these highly procedural, context-dependent, and contestable values. This helps the designers evade the formalism trap [Selbst et al., 2019].

Value Dams and Flows are used for making decisions regarding value tensions between stakeholders. Value tensions occur when the values of stakeholders conflict with one another. In this method, value dams represent the technical features or organizational policies that some stakeholders strongly oppose. The main idea is that strong opposition to a feature or policy can inhibit effective use of the technology, or even block appropriation of the technology entirely. Therefore it is very important to identify these oppositions by the use of empirical methods such as interviews and surveys. Value flows are the direct opposite of value dams. They represent the features and policies which a large number of stakeholders are in favor of incorporating. The main idea is that implementing value flows in the design of the technical tools or organizational policies may attract stakeholders to adopt the tool [Miller et al., 2007].

5.1.2 Participatory Design

Participatory Design, also referred to as co-creation, is a human-centric design approach where the key stakeholders of a system are involved in its design process. This is due to participation and democracy being the core values of Participatory Design [van den Hoven et al., 2015]. These methods directly contradict traditional design approaches, which are based on a distance between the designers and future users of a system [Toni Robertson, 2012]. This leads to a better understanding of the needs and values of the stakeholders [Elizarova et al., 2017], exploring alternative visions, and giving otherwise "invisible" individuals a voice [van der Velden and Mortberg, 2015].

Because of these reasons, Participatory Design is a very important methodology when designing for values embodied by human rights, such as autonomy over self-representation.

Additionally, Participatory Design brings various benefits with it, such as the ability to avoid the solutionism trap—by exploring alternative visions for the proposed solution—and the formalism trap by getting a clearer view of the implicated values [Selbst et al., 2019].

There are a wide range of methods available to facilitate inclusive and cooperative design, such as storyboards [Babich, 2017], ethnography [Crabtree, 1998], collaborative prototyping, etc.

Ethnography is particularly used in the initial phases of the design process, to get a deeper understanding of the social context and its stakeholders; and to elicit requirements the system should possess according to the needs of the stakeholders. The ethnographer may also double as a facilitator in other Participatory Design methods, to create meaningful dialogues and enable collaboration [Blomberg et al., 1993]. These collaborations try to close the distance between the various stakeholders and let everyone contribute with their knowledge and views.

The next subsection—Systemic Design—explores another STSD approach that complements the methods covered in Design for Values.

5.2 Systemic Design

Design has evolved from a practice aimed at designing products and services, into one capable of addressing complex societal challenges. This is due to a shift towards social innovation. Social innovation is the practice of designing solutions for societal problems in such a way that the value created by the solution benefits society as a whole rather than the private individual [Phills Jr. et al., 2008]. This practice contradicts traditional product design by expanding the focus on multiple fronts, such as expanding from users and customers to society as a whole; and from designing products and services to designing complex policies, strategies, and social systems. This shift from traditional product design to social innovation means traditional design practices have to adapt to this new context. One of these adaptations is a shift towards systems thinking.

Capra [1997] describes a system as "an integrated whole whose essential properties arise from the relationships between its parts". None of the parts have these essential properties separately. Systems thinking therefore considers everything in relation to a larger system/indivisible whole of which they are part; this philosophy is called synthesis. This contradicts reductionist approaches to complex problem solving, which view a system as a whole that should be taken apart. However, Ackoff [2004] argues that improving the parts of a system separately usually does not lead to an improvement of the system as a whole.

Systemic Design is a unified field of practice combining elements of systems thinking with elements of design, in the context of social innovation. This design practice builds on the analytical benefits of systems thinking and the action-oriented benefits of design [Buchanan, 2019].

Systemic Design is highlighted in this paper due to the broad view on any social system it offers. Specifically, since Systemic Design combines systems thinking with design practices, one could combine the broad mindset of systems thinking with design methodologies from Design for Values (section 5.1). These two STSD approaches complement each other and combining them results in a better view of the social system.

van der Bijl-Brouwer and Malcolm [2020] analyzed reoccurring Systemic Design principles underlying the practices of different schools of thought within Systemic Design. These design principles provide a base when designing a socio-technical system and are combined with other STSD approaches.

A Systemic Design principle that complements Design for Values concerns opening up the problem space and acknowledging the interrelatedness of problems. This principle combines systemic thinking with the problem framing methodology from design, leading to the problem space being expanded. Specifically, due to the perspective taken in systems thinking, a more comprehensive understanding can be obtained by expanding the system rather than reducing it [Ackoff, 1999]. Using this viewpoint when framing the problem leads to different perspectives on the problem and opens up the design space to more possible solutions. This design principle helps evade the framing trap—which arises by failing to model the entire system correctly due to a too narrow scope [Selbst et al., 2019]—which Design for Values is not equipped to evade by itself. Another trap this design principle helps evade is the solutionism trap, which arises by exclusively focusing on developing a technological solution [Selbst et al., 2019]. There are various concrete tools available in systems thinking to generate the expanded system, such as the rich picture tool [Checkland, 1981] and concept mapping [Ltd., 2019].

Another design area that combines well with systems thinking is human-centered design (HCD). HCD focuses on a deep understanding of human beings and involves human perspective throughout the design process to address the human aspects of a social system [Bijl-Brouwer and Dorst, 2017]. Systemic Designs' focus on human relationships, coupled with HCD's understanding of human beings, results in better exploration of existing relationships, value tensions, and human-centered solutions.

This design principle can be used when engaging stakeholders, for instance. Instead of exclusively focusing on the needs of each individual stakeholder, one should focus on the bigger picture, namely the relationship between stakeholders. This mindset should be used when applying stakeholder engagement methods from Participatory Design, Requirements Engineering, or Value Sensitive Design.

Finally, designers should follow an iterative design methodology throughout the design process when addressing complex problems. The relevant design area here is co-evolution, and it follows an experimental and evolutionary approach to problem-solving. By making use of prototypes [Gadiraju and Tielma, 2021b], designers can both safely test their understanding of the problem and test the reaction of key stakeholders to their current solution [Dorst and Cross, 2001].

This section described the motivation for using STSD, followed by a description of Design for Values, its building blocks Value Sensitive Design and Participatory Design, and various methods one should consider when designing a socio-technical solution. Additionally, I have explained how Systemic Design can complement Design for Values by avoiding the framing and solutionism traps, focusing on relationships between stakeholders, and following an iterative design approach.

6 Responsible Research

This paper ensures proper citation and academic honesty in the use of its references. Additionally, no work has been excluded due to being published by a specific person, organisation, or having a controversial/contrasting theory. Therefore the discussed papers have been selected in a fair way. However, although the information can be fact-checked—like the definitions, advantages, and disadvantages of certain methods—one might disagree with the drawn conclusions and the way I bridge this socio-technical gap. One might prefer other STSD approaches, or not even agree STSD approaches are the way to go in the first place.

7 Discussion

This section starts off by summarizing the issue at hand, the research question, and the main findings. This is followed by initial thought on how the designers should start the design process. Next, an analysis of the multidisciplinary design team involved in the STSD is presented. Finally, the limitations and directions for future research are discussed.

Summary The main issue discussed in this paper is that AI-powered tools used in hiring practices can violate the applicants' autonomy over self-representation. This is because generally, machine learning algorithms assign numerical values to a set amount of target variables. However, the target variables present in hiring practices are context-dependent, and thus contestable [Selbst et al., 2019]. Therefore when the algorithm tries to numerically quantify these complex variables, it ends up stripping away the context from these variables [Aizenberg and van den Hoven, 2020], which results in the applicant being misrepresented [Delandshere and Petrosky, 1998]. The applicant's autonomy over self-representation is infringed when the algorithm's conception of the applicant becomes the focus of the hiring organization since the applicant is being misrepresented and they are not in control of their own representation anymore [Broeders et al., 2017].

Designers should consider that purely focusing on the technical features of these machine learning algorithms might not be enough to ensure the applicants' autonomy over self-representation [Selbst et al., 2019]. There are additional factors one should consider, such as human, social, and organizational factors. Therefore this paper proposes a broader design scope called "Socio-Technical Systems Design (STSD)."

The main research question is *How can socio-technical systems design approaches ensure autonomy in hiring practices?* with as a goal to explore and pick suitable STSD approaches which can be applied in hiring practices to ensure the applicants' autonomy over self-representation. Two complementing STSD approaches have been selected after conducting a literature study: Design for Values and Systemic Design.

Designers have to get a nuanced understanding of the social context, its politics, and stakeholders, before developing a solution. This paper suggests that getting this understanding can be gathered by involving key stakeholders in the design process [Elizarova et al., 2017, van den Hoven et al., 2015]. Therefore the designers should initially conduct a stakeholder analysis to identify the key stakeholders in the social context [Friedman et al., 2006, van den Hoven et al., 2015], which should essentially include stakeholders from the hiring organization and the applicant.

After identifying and assembling the key stakeholders an empirical investigation should be conducted to explore the stakeholders' needs and the values which the proposed system might implicate. There are a wide range of methods available to elicit these requirements, spanning various design methodologies such as Participatory Design [van der Velden and Mortberg, 2015], Human-Centered Design [Bijl-Brouwer and Dorst, 2017], and Requirements Engineering [Zowghi and Coulin, 2005].

The mindset offered by Systemic Design helps in broadening up the design scope as it considers everything in relation to a larger indivisible whole [van der Bijl-Brouwer and Malcolm, 2020]. This brings multiple benefits with it, like that different perspectives on any given problem, in turn, make place for more possible solutions. Additionally, by following an iterative design approach the designers can safely test their understanding of the problem and the reaction of key stakeholders to their current solution [Dorst and Cross, 2001, Gadiraju and Tielma, 2021b].

This nuanced, stakeholder-centric approach results in an inclusive, transparent, multi-discipline, and socially aware process which is necessary to understand the complex social context, and its socio-ethical issues. The designers should work together with the key stakeholders to get answers on topics like: *What is important for each stakeholder? Are there value* *tensions between the stakeholders? Which design requirements should a proposed solution have to respect the stakeholders' values? What makes a good candidate? What is competence? How can competence be assessed and measured? Which aspects can be quantified? Does measurement x cover the candidates' universal domain or is the measured variable more complex than that?*

This list is not exhaustive, but asking these kinds of questions will help designers get a deeper understanding of the social context, the stakeholders' values, and how they should design the system.

Initially, the designers could conduct interviews with recruiters and talent acquisition specialists to find out which tools they are using at the moment to see how a new solution could fit into this process. Additionally, the designers could explore what makes a candidate good and which competencies are important in a potential employee by working together with current employees and project/team leaders. After knowing which competencies are important in a prospective employee, the designers could conduct research to find out what the best way is to measure these competencies. Some competencies may be quantifiable, others might require a qualitative assessment. Also, designers could explore the applicant's perspective through story-boarding [Babich, 2017] or scenarios [Zowghi and Coulin, 2005], for instance. After accumulating the values of the applicants and the stakeholders involved in the hiring organizations, the designers can translate these values into specific design requirements by specifying them in a value hierarchy [van de Poel, 2013], for instance, after which a focused debate between the stakeholders can be held. Moreover, value tensions can be identified and mitigated at this stage [Miller et al., 2007]. Finally, by working iteratively and making use of prototypes the designers can always check the reception of their proposed solution with the key stakeholders [Gadiraju and Tielma, 2021b].

The Design Team This paper referred to "the designers" of these systems without specifying who these designers are. Developing a socially conscious solution for this context requires multiple disciplines to work together. Three disciplines are covered here due to their complementing roles in the design process: social science, computer science, and design.

Social scientists aim at making informed and reasoned decisions for the public good [Lewison, 2020]. It plays an incredibly important role in designing a fair and human-centric solution to the problem at hand. The education of a social scientist has an emphasis on human behavior and relationships. Therefore, a social scientist, such as an ethnographer (section 5.1.2), is better equipped to perform the stakeholder analysis and engagement techniques than a computer scientist, for instance. To be more specific, a social scientist would be a great contribution to a development team by eliciting requirements in the form of interviewing stakeholders, facilitating sessions (such as focus groups), etc.

Design specialists are necessary throughout the development cycle to embed the stakeholders' values in the product. This ranges from exploring the implicated values through prototypes [Gadiraju and Tielma, 2021b], to actually embedding the values in the user interface by a UX/UI designer [ter Haar Romenij, 2020]. Additionally, designers are knowledgeable in how to perform expert/user evaluations, controlled experiments, etc. These abilities are necessary to make sure that the product stays in line with the needs of the stakeholders'.

Computer scientists are generally educated to develop software and analytical skills. Therefore they are well equipped to perform logical and investigative tasks, such as developing software, designing the information architecture of a system [Gadiraju and Tielma, 2021a], etc.

However, computer scientists have a big responsibility when developing software. They should be aware and mindful of the socio-technical risks AI applications bring with them, such as discrimination [Shin, 2020] and data determinism (section 3). Therefore it is important to educate computer scientists about the ethics and possible dangers of these applications. An example of this is a course given at the TU Delft called "TI0388WM - IT and Values", which had as a goal to teach students how to "critically assess ethical, societal, political, and epistemological arguments related to IT" (Duran, 2019, p. 1). By making computer scientists think about the ethical considerations of AI during their education, they will be better equipped to actively consider whether the application they are developing is being done in an ethically sound way.

This is a complex social context that demands a broad design scope and multiple disciplines working together; there is not one sole expert which has all of the answers. Therefore when opting for a socially conscious system, these disciplines should collaborate and complement each other's competencies. However, this multidisciplinary process has its limitations. These limitations are mainly due to the borders between the different disciplines [Baxter and Sommerville, 2011]. Bader and Nyce [1998] explain that this is due to the disciplines not fully understanding what the other disciplines can offer, leading to none of the disciplines asking for assistance in the development process.

This issue can be tackled by exposing disciplines to each other, either during education or through practice. By familiarizing different disciplines they may have an easier time working together by understanding each others' possible contributions. For instance, computer scientists can be given a course about STSD during their education. The course could be given by a designer and/or social scientist. By exposing computer scientists to other disciplines early on, they could be better equipped to understand the possible contribution of each discipline. Additionally, teaching computer scientists about STSD shows them that there are alternative design approaches that hold the values of stakeholders central, which may help them evade the solutionism trap in the future [Selbst et al., 2019].

Limitations This paper is theoretical. The information presented in this paper has been obtained by conducting a literature study, applying it to hiring practices, and drawing conclusions from it. This means that although this paper references papers where empirical research has been conducted, it does not conduct any empirical research itself. Therefore, it is not confirmed whether combining these STSD approaches actually works in practice. Additionally, due to the process of designing for values being context-dependent, and therefore contestable, this paper can not ensure that the presented STSD approaches are optimal. Therefore one might find another STSD approach that is more suitable for ensuring the applicants' autonomy in the selection process.

Future Work Future research can build on these limitations by applying STSD approaches in practice. More specifically, that means exploring various stakeholder analysis techniques to identify the key stakeholders in hiring practices, followed by engaging with the stakeholders and explore their needs and values. Researchers can explore the advantages and disadvantages of different engagement techniques by actually applying different methodologies in practice. This also counts for STSD approaches. By normalizing the use of STSD in practice, more research and knowledge can be built on it; leading to refinement of existing methodologies and the creation of new ones.

References

R. L. Ackoff. *Ackoff's best: His classic writings on management*. New York: Wiley, 1999.

R. L. Ackoff. Transofrming the systems movement, 2004. URL https://thesystemsthinker.com/wp-content/uploads/pdfs/150801pk.pdf.

AI Myths. *The term AI has a clear meaning*. n.d. URL https://www.aimyths.org/the-term-ai-has-a-clear-meaning#alternative-terms. Publication Title: aimyths.

E. Aizenberg and J. van den Hoven. Designing for human rights in AI. *Big Data & Society*, 7 (2):2053951720949566, July 2020. ISSN 2053-9517. doi: 10.1177/2053951720949566. URL https://doi.org/10.1177/2053951720949566. Publisher: SAGE Publications Ltd.

E. T. Albert. AI in talent acquisition: a review of AI-applications used in recruitment and selection. *Strategic HR Review*, 18(5):215–221, Oct. 2019. ISSN 1475-4398, 1475-4398. doi: 10.1108/SHR-04-2019-0024. URL https://www.emerald.com/insight/content/doi/10.1108/SHR-04-2019-0024/full/html.

N. Babich, Oct 2017. URL https://www.smashingmagazine.com/2017/10/storyboarding-ux-design/.

G. Bader and J. M. Nyce. When only the self is real: Theory and practice in the development community. *Journal of Computer Documentation*, 22(1):5–10, 1998. ISSN 0731-1001. URL https://www.learntechlib.org/p/84220.

G. Baxter and I. Sommerville. Socio-technical systems: From design methods to systems engineering. *Interacting with Computers*, 23(1):4–17, Jan. 2011. ISSN 09535438. doi: 10.1016/j.intcom.2010.07.003. URL https://academic.oup.com/iwc/article-lookup/doi/10.1016/j.intcom.2010.07.003.

P. M. Bednar and C. Welch. Socio-Technical Perspectives on Smart Working: Creating Meaningful and Sustainable Systems. *Information Systems Frontiers*, 22(2):281–298, Apr. 2020. ISSN 1387-3326, 1572-9419. doi: 10.1007/s10796-019-09921-1. URL http://link.springer.com/10.1007/s10796-019-09921-1.

M. v. d. Bijl-Brouwer and K. Dorst. Advancing the strategic impact of human-centred design. *Design Studies*, 53:1–23, 2017. ISSN 0142-694X. doi: https://doi.org/10.1016/j.destud.2017.06.003. URL https://www.sciencedirect.com/science/article/pii/S0142694X17300455.

J. Blomberg, J. Giacomi, A. Mosher, and P. Swenton-Wall. Ethnographic field methods and their relation to design. In D. Schuer and A. Makioka, editors, *Participatory design: Principles and practices*, pages 123–156+. Erlbaum, 1993.

D. Broeders, E. Schrijvers, and E. H. Ballin. Big data and security policies: Serving security, protecting freedom. page 34, 2017.

R. Buchanan. Systems Thinking and Design Thinking: The Search for Principles in the World We Are Making. *She Ji: The Journal of Design, Economics, and Innovation*, 5 (2):85–104, 2019. ISSN 2405-8726. doi: https://doi.org/10.1016/j.sheji.2019.04.001. URL https://www.sciencedirect.com/science/article/pii/S2405872618301370.

T. Caestano and B. Simpson-Young. Artificial intelligence can deepen social inequality. Here are 5 ways to help prevent this, 2021. URL https://theconversation.com/artificial-intelligence-can-deepen-social-inequality-here-are-5-ways-to-help-prevent-this-15222

F. Capra. *The web of life: A new synthesis of mind and matter*. 1997.

P. Checkland. *Systems Thinking, Systems Practice*. 1981.

A. N. Christakis and K. C. Bausch. How people harness their collective wisdom and power to construct the future in co-laboratories of democracy, 2006.

C. W. Clegg. Sociotechnical principles for system design, Aug 2000. URL https://www.sciencedirect.com/science/article/abs/pii/S0003687000000090?via\%3Dihub.

A. Crabtree. Ethnography in Participatory Design. page 14, 1998.

S. Daley. 32 examples of AI in healthcare that will make you feel better about the future. *Builtin*, Apr. 2019. URL https://builtin.com/artificial-intelligence/artificial-intelligence-healthcare.

G. Delandshere and A. Petrosky. Assessment of Complex Performances: Limitations of Key Measurement Assumptions. page 12, 1998.

C. Detweiler and M. Harbers. Value Stories: Putting Human Values into Requirements Engineering. page 10, 2014.

K. Dorst and N. Cross. Creativity in the design process: co-evolution of problem-solution. *Design Studies*, 22(5):425–437, 2001. ISSN 0142-694X. doi: https://doi.org/10.1016/S0142-694X(01)00009-6. URL https://www.sciencedirect.com/science/article/pii/S0142694X01000096.

J. Dryden. Autonomy. URL https://iep.utm.edu/autonomy/.

J. M. Duran. Course browser searcher, 2019. URL https://studiegids.tudelft.nl/a101_displayCourse.do?course_id=52247.

O. Elizarova, J. Briselli, and K. Dowd. Participatory design in practice, 2017. URL https://uxmag.com/articles/participatory-design-in-practice.

F. Emery and E. Trist. Socio-technical systems, 1960.

J. Fernando. Learn what stakeholders are and the roles that they play, May 2021. URL https://www.investopedia.com/terms/s/stakeholder.asp.

B. Friedman. *Human Values and the Design of Computer Technology*. 1997.

B. Friedman and D. Hendry. The envisioning cards: A toolkit for catalyzing humanistic and technical imaginations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pages 1145–1148, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450310154. doi: 10.1145/2207676.2208562. URL https://doi.org/10.1145/2207676.2208562.

B. Friedman, P. Kahn, and A. Borning. *Human-computer interaction in management information systems: foundations*. 2006.

U. Gadiraju and M. Tielma. Lecture 2 notes in human computer interaction, May 2021a.

U. Gadiraju and M. Tielma. Lecture 4 notes in human computer interaction, May 2021b.

S. Gudavajhala. Top 5 stakeholder analysis techniques in projects: Stakeholder analysis in project management, 2017. URL https://masterofproject.com/blog/7514/top-5-stakeholders-analysis-techniques-in-projects-stakeholder-analysis-in-project-management.

M. Halbertal. Three concepts of human dignity, Nov. 2015. URL https://youtu.be/FyEvREFZVvc.

C. Kerry. Protecting privacy in an AI-driven world, 2020. URL https://www.brookings.edu/research/protecting-privacy-in-an-ai-driven-world/.

R. Kitchin. Big Data, new epistemologies and paradigm shifts. *Big Data & Society*, 1(1): 205395171452848, Apr. 2014. ISSN 2053-9517, 2053-9517. doi: 10.1177/2053951714528481. URL http://journals.sagepub.com/doi/10.1177/2053951714528481.

D. Levin. Professional Associations & Organizations, 2016. URL https://jobstars.com/professional-associations-organizations/.

K. Lewison. Why is social studies important? 8 reasons to study, May 2020. URL https://www.uopeople.edu/blog/why-is-social-studies-important/.

X. Ltd. What is a concept map and how to do concept mapping (actionable guide for beginners), 2019. URL https://www.xmind.net/blog/en/concept-map-tutorial/.

S. C. Matz, M. Kosinski, G. Nave, and D. J. Stillwell. Psychological targeting as an effective approach to digital mass persuasion. *Proceedings of the National Academy of Sciences*, 114 (48):12714–12719, Nov. 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1710966114. URL http://www.pnas.org/lookup/doi/10.1073/pnas.1710966114.

J. Miller, B. Friedman, and G. Jancke. Value tensions in design: The value sensitive design, development, and appropriation of a corporation's. pages 281–290, 01 2007. doi: 10.1145/1316624.1316668.

L. P. Nathan, B. Friedman, P. Klasnja, S. K. Kane, and J. K. Miller. Envisioning systemic effects on persons and society throughout interactive system design. In *Proceedings of the 7th ACM Conference on Designing Interactive Systems*, DIS '08, pages 1–10, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605580029. doi: 10.1145/1394445.1394446. URL https://doi.org/10.1145/1394445.1394446.

Official Journal of the European Union. Charter of fundamental rights of the european union, 2012.

J. A. Phills Jr., K. Deiglmeier, and D. T. Miller. Rediscovering social innovation (ssir), 2008. URL https://ssir.org/articles/entry/rediscovering_social_innovation#.

Sas. Artificial Intelligence - What it is and why it matters, 2021. URL https://www.sas.com/en_us/insights/analytics/what-is-artificial-intelligence.html.

D. Schuler and A. Namioka. *Participatory Design: Principles and Practices*. 1993.

A. D. Selbst, D. Boyd, S. A. Friedler, S. Venkatasubramanian, and J. Vertesi. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, pages 59–68, New York, NY, USA, Jan. 2019. Association for Computing Machinery. ISBN 978-1-4503-6125-5. doi: 10.1145/3287560.3287598. URL https://doi.org/10.1145/3287560.3287598.

O. Sensen, editor. *Kant on Moral Autonomy*. Cambridge University Press, 2012. doi: 10.1017/CBO9780511792489.

T. Shin. Real-life Examples of Discriminating Artificial Intelligence, 2020. URL https://towardsdatascience.com/real-life-examples-of-discriminating-artificial-intelligence-cae395a90070.

M. Singh. Objectives of Recruitment and Selection, 2020. URL https://blog.mettl.com/guide/objectives-of-recruitment-and-selection/.

Steven. Who are the main recruitment stakeholders in the recruitment process?, Apr 2020. URL https://www.recruitive.com/who-are-the-key-recruitment-stakeholders/.

J. ter Haar Romenij. Empowering Academic Graduate Job Search, 2020.

I. W. Toni Robertson. Ethics. In *Routledge International Handbook of Participatory Design*. Routledge, 2012. doi: 10.4324/9780203108543.ch4. URL https:. Section: chapter4.

University of Siegen and F. Yetim. Bringing Discourse Ethics to Value Sensitive Design: Pathways toward a Deliberative Future. *AIS Transactions on Human-Computer Interaction*, 3(2):133–155, June 2011. ISSN 19443900. doi: 10.17705/1thci.00030. URL https://aisel.aisnet.org/thci/vol3/iss2/5.

I. van de Poel. Translating Values into Design Requirements. In D. P. Michelfelder, N. McCarthy, and D. E. Goldberg, editors, *Philosophy and Engineering: Reflections on Practice, Principles and Process*, pages 253–266. Springer Netherlands, Dordrecht, 2013. ISBN 978-94-007-7762-0. doi: 10.1007/978-94-007-7762-0_20. URL https://doi.org/10.1007/978-94-007-7762-0_20.

J. van den Hoven, P. E. Vermaas, and I. van de Poel, editors. *Handbook of Ethics, Values, and Technological Design*. Springer Netherlands, Dordrecht, 2015. ISBN 978-94-007-6969-4 978-94-007-6970-0. doi: 10.1007/978-94-007-6970-0. URL http://link.springer.com/10.1007/978-94-007-6970-0.

M. van der Bijl-Brouwer and B. Malcolm. Systemic Design Principles in Social Innovation: A Study of Expert Practices and Design Rationales. *She Ji: The Journal of Design, Economics, and Innovation*, 6(3):386–407, 2020. ISSN 24058726. doi: 10.1016/j.sheji.2020.06.001. URL https://linkinghub.elsevier.com/retrieve/pii/S2405872620300447.

M. van der Velden and C. Mortberg. Participatory design and design for values. In J. van den Hoven, P. E. Vermaas, and I. van de Poel, editors, *Handbook of Ethics, Values, and Technological Design*, chapter 1. Springer Netherlands, 2015. ISBN 978-94-007-6969-4 978-94-007-6970-0. doi: 10.1007/978-94-007-6970-0. URL http://link.springer.com/10.1007/978-94-007-6970-0.

A. Westin. *Privacy and Freedom*. 1970.

E. Yong. A Popular Algorithm Is No Better at Predicting Crimes Than Random People. *The Atlantic*, Jan. 2018. URL https://www.theatlantic.com/technology/archive/2018/01/equivant-compas-algorithm/550646/.

D. Zowghi and C. Coulin. Requirements Elicitation: A Survey of Techniques, Approaches, and Tools. In A. Aurum and C. Wohlin, editors, *Engineering and Managing Software Requirements*, pages 19–46. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. ISBN 978-3-540-28244-0. doi: 10.1007/3-540-28244-0_2. URL https://doi.org/10.1007/3-540-28244-0_2.