Delft University of Technology

CSE3000 Research Project

Using Machine Learning for University Admission: Mapping the Socio-Technical Issue

June 2021

Author: Omri Niri

Supervisor: Dr. Evgeni Aizenberg

Responsible professor: Prof. Dr. Ir. R.L. (Inald) Lagendijk

Contents

Abstract

Machine learning algorithms were used in in the past decade to assist humans with recruitment and grades assessments in the academic field. For the most part, the algorithms either exacerbated existing biases or output unfair results. This could often be traced back to an ill-implementation of the systems in the social context. The academic admission process is defined as setting goals, locating candidates, ranking and accepting them. To properly integrate machine learning in such process, one may follow the Value Sensitive methodology, which suggests designing a technical system around a social value. This methodology takes into account the various stakeholders, values and technical solutions available. Later, the system should be iteratively improved and constantly evaluated and examined so that it still serves the core values as defined.

1) Introduction

Recent advancements in hardware, software and algorithms allow us to use computers in ways like never before. The computer's immense speed and ability to handle complex problems make it tempting for humans to allow algorithms increasing degrees of autonomy in many decision-making processes in our society. Specifically, a certain group of algorithms has been made more accessible and easier to use in various contexts: different forms of machine learning algorithms (ML), sometimes also referred to as primitive artificial intelligence (AI), now have simple application programming interfaces (APIs), which make them easier to integrate in new software. Their popularity increases and we see ML used more and more in new areas.[1] ML is the collective name given to *"computer algorithms that improve automatically through experience and by the use of data"*.[2] ML systems modify their internal model of the problem automatically according to pre-defined principles. A critical aspect when using this family of algorithms, is that it could be difficult to backtrack what lead to a certain outcome.[3]

While the term 'ML' could be used to describe a specific algorithm (e.g., linear regression, K-nearest neighbors, a combination of such, etc.), the term 'AI' is broader and has more extensive connotations in today's lexicon. For the purposes of this paper, I shall often use the terms ML and AI interchangeably and the meaning will be clear from the context. Also, acknowledging the quick pace at which AI is progressing, whenever I mention AI, it refers to the state-of-the-art AI systems as they exist in the year 2021.

This paper is a qualitative examination of the academic admission process, and how AI could be used in it. I shall review the field as it is today including some case studies, and make suggestions for it going forward, borrowing from practices in related fields. In current literature, attempts to explicitly characterize AI-supported academic admissions are still in their infancy. Most of the existing research focuses on predicting the admission decision by a given institution using ML,[4,5] or the prediction of pupils' future performance given their current performance.[6] While there have been quite a few papers written about the existence and effects of bias in AI in general,[7–11] it is hard to find any studies that relate directly to academic admission. Useful articles can be naturally borrowed, mutatis mutandis, from studies about candidates screening in the job market, which are more elaborate. Such work is that of Liem et al.,[12] Zehlike et al.[13] and Derous et al.[14], to name a few.

Although studies analyzing the bias in ML-supported admission are quite scarce, the topic is in dire need of examination. Recent real-world use of ML-supported admission processes encountered difficulties and unexpected outcomes, where the new system reenforced existing bias instead of decreasing it. This example and others will be discussed further in chapter 3.

Before moving forward, there are two key concepts the reader should be familiar with.

The first concept is the socio-technical perspective. In its core, it is the realization that any technical artifacts are embedded within a social context, and that it is less meaningful to take into account only the technical aspects of a system when designing it. Designs should not only meet technical requirements (e.g., speed, database type, menu format etc.) but also attend to human-machine interaction and adhere to social calls.[15] Implications of these two seemingly unrelated approaches brought together, in the context of ML-enhanced academic admission systems, will be discussed in depth throughout this paper.

The second crucial concept is the so-called solutionism trap. Selbst et al. define the *solutionism trap* as *"Failure to recognize the possibility that the best solution to a problem may not involve technology"*.[16] This is a major point, crucial to this discussion. Engineers tend to turn to technology automatically, not considering the fact that an appropriate solution might come from a non-technical discipline. Sometimes it is a social issue, which cannot or should not be mitigated using any innovative technology.

Being familiar with the two concepts above allows a more fluent discussion about the admission process, which is being defined in the next chapter. From as early as when the institution's goals are defined until some candidates are accepted, chapter 2 breaks down and gives names to the different steps of the admission process. Then, chapter 3 is dedicated to one particular step of the process, namely- the step of ranking the candidates. It discusses the relation between the institution's goals and the candidates' properties. In chapter 4 I introduce two common socio-technical methodologies, which are natural candidates for design approach while integrating ML into the admission process. Chapter 5 evaluates each method in its context and determines which is more suitable. Then, chapter 6 progresses to the actual application. A discussion and further recommendations are given in chapter 7.

2) The admission process: definitions

In an academic context, the admission process may be applied in various cases and in various manners. A rather trivial situation is a study program which has no maximal capacity for the students it may accommodate. This can happen in a Bachelor degree program, for example, accepting all candidates who meet or exceed specific requirements. There is no limit on the number of students that the program can accept and no prioritization of the candidates is needed, as every individual will be considered regardless of their peers.

A less trivial admission process takes place in certain academic programs, from Bachelor and Master degrees to research positions one needs to fulfill, where the institution cannot admit an unlimited number of candidates. In this case, the institution usually conducts a *comparative* admission process. Let us shortly and generically break down such a process into four steps, which I define as follows:

1. The academic institution defines its *goals*. This is done regardless of any admission process, and it is used a "moral compass" for the organization. These goals could be explicit and specific or implicit and vague, as later examples will show when I expand on this step.
2. The institution *locates candidates* for its vacant positions, be it students or researchers.

3. The institution *ranks the candidates* based on their expected contribution to its defined goals from step 1. This is not to say that the individuals only benefit the organization and not vice-versa, but rather to emphasize that the institution prefers the candidates that contribute to its own goals.
4. The institution *accepts*, perhaps after a negotiation, the top candidates until all vacancies are fulfilled.

While step 4 is necessary mostly from an administrative point of view, steps 1-3 have more to do with fairness and socio-technical issues. This paper's focus is on steps 1-3 and I shall now elaborate about important aspects of each. Later, I will discuss the design of the process from a socio-technical perspective.

Step 1: Defining the institution's goals

What an ideal candidate is, depends greatly on the institution's goals. Those are usually determined regardless of any admission process the institution conducts. Goals serve multiple purposes in the organization's view, like establishing priorities and defining the relationship between the institution and society.[17] From short-term to long-term, a desired candidate should help the institution achieve these goals while probably also benefiting from the organization of course, materialistically or otherwise, like in any symbiotic relationship. This point of view of the applicants will be central for later analysis.

Defining goals is the step where an organization may express its purpose as a list of goals derived from values. For example, Delft University of Technology, The Netherlands, defines its vision with a rather broad declaration: *"Delft University of Technology contributes to solving global challenges by educating new generations of socially responsible engineers and expanding the frontiers of the engineering sciences."*[18] A vague, broad declaration indeed. The reader might be left wondering what these global challenges are, what does "responsible" mean in this context and more.

It is worth reviewing what the most common goals amongst universities are. A survey done by Gross et al.[19] shows which goals members of academic institutions perceive as most important. They interviewed members of different universities, both academics as well as administrators, and found that there seems to be a consensus between the two types of personnel. Out of forty-seven suggested goals, the top four are a mixture between academic and administrative goals: (1) protecting academic freedom, (2) increasing the university's prestige, (3) ensuring donations, and (4) training scholars.

It is not a coincidence that many organizations and the individuals working there opt to define their goals in a vague manner. Formulating explicit and exhaustive goals for an organization is hard, especially due to the existence of the institution's many stakeholders, each with their own interests.[17] This complex web of connections will be explored in chapter 5, under stakeholders analysis.

Step 2: Locating candidates

The second step of the admission process is to assemble a set of candidates for the vacant position(s), a list which will be sorted in step 3. Although the task of finding candidates may sound like a trivial technical step, social considerations apply here as well. Let us illustrate these considerations with a hypothetic example.

Consider a vacant research position at a certain academic institution, for which less applications than expected are received. The institution decides to seek out and attract more candidates to increase its selection, so it initiates an advertisement campaign on social media for the position. Since men and women have different usage patterns of social networks[20], the final candidate list is likely to be somewhat biased for or against a certain gender. Men-women usage ratio varies between social networks, which strengthens the need to examine where one is looking for candidates.

This argument applies when the institution performs a passive search (i.e., trusting the university's reputation or advertising to attract candidates) as well as when performing an active search (i.e., reaching out to specific individuals and making them an offer). Do note that these considerations are not directly or necessarily related to the use of ML in this step, but in general.

Step 3: Ranking candidates

Once a set of potential candidates is established, the ranking process can take place. In this paper, I only consider the case in which the number of candidates is strictly larger than the number of vacant positions and the selection is not random. The key question of this step is therefore, how to rank the candidates? The ranking step itself consists of at least two distinguishable but related sub-steps, which are at the heart of the socio-technical challenge for academic admission: Choosing properties by which to sort, and determining the ranking algorithm (i.e., the weight given to each property and the ranking method itself). As this is a core matter in the admission process, I dedicate the entirety of the next chapter to this step.

3) Important considerations in ranking candidates

In order to choose which properties to rank its candidates by, an institution may want to refer back to its own goals, as they were defined in step 1. From the institution's perspective, a chosen candidate should contribute to the institution's goals. So, the main aim of the ranking step should be finding those candidates which, based on their properties, are estimated to contribute to the organization's goals the most. Making this kind of prediction is hard, as we are trying to correlate abstract ideas with each other, such as generic goals and human properties.

In natural sciences and mathematics, finding correlation between variables is often a matter of mathematical formulas. Common techniques such as linear regression or r-score can be used when the given variables are measurable and well defined. However, in order to apply such techniques to correlate the institution's goals with the candidates' properties, the challenge is on two fronts: First, the goals themselves are often vague, as shown in the previous chapter. Second, while some properties come already in the form of measurable numbers (e.g., high school Grade Point Average (GPA), Scholastic Assessment Test score (SAT) etc.) other properties are not always easily quantifiable, if they can be quantified at all (e.g., leadership skills, teamwork abilities, ambition etc.).

Yet, the contention used to justify any candidates ranking is that it is possible to predict future behavior and performance based on existing data. In this case, this translates into predicting how the organization's vague goals will be best fulfilled by a candidate which possesses certain quantifiable and unquantifiable properties. If an institution seeks to make such ambitious prediction, especially using ML, it has to bridge a socio-technical gap between the social context (of the institution's goals, the candidates' properties and the purpose of the whole admission process) and the necessary technical formality of algorithms.

I will now lay out several case studies, to see what some universities claim to be looking for in academic candidates for undergraduate programs. For its admission process, Harvard College, USA, says that *"there is no such thing as a typical Harvard student"*. Instead, they claim to manually process applications while giving intuitive thought to broadly defined properties, such as *"growth and potential"*, *"interests and* *activities"*, *"personal character"*, *"(potential) contribution to the Harvard community"* and more.[21] In 2009, the dean of Harvard, W.R. Fitzsimmons, said the following about their admission process:

> *"While we value objective criteria, we apply a more expansive view... Test scores and grades offer some indication of students' academic promise and achievement. But we also scrutinize applications for extracurricular distinction and personal qualities…*
>
> *Efforts to define and identify precise elements of character, and to determine how much weight they should be given in the admissions process, require discretion and judiciousness. But the committee believes that the 'best' freshman class is more likely to result if we bring evaluation of character and personality into decisions than if we do not...*
>
> *These considerations are guidelines that are neither comprehensive nor absolute."*[22]

A similar approach is displayed by other top-ranking universities in the USA, such as Stanford, MIT and Yale.[23–25] These are amongst the top ranking universities in the world, with ample resources allowing for a meticulous admission process and a wide selection of applicants, and I use these institutions as the paper's benchmark. It is important to keep this context in mind. This small selection of universities is what I use here, even though they are of a very specific kind: Top universities in the United States. Other institutions in other countries, or even within the USA, may have different resources, face different circumstances upon their admission processes and prefer different admission practices.

As shown, the benchmark universities, which have more resources available to them than to most academic institutions in the world, use human judgement for the most crucial part of the admission process, the evaluation itself. This leads to the following intermediate conclusion: As of 2021 and according to the benchmark-universities' practices, the quality of human-based admission still outweighs the speed and efficiency of AI, at least in the eyes of leading universities in the USA.

At least one concrete justification for this choice, to incorporate humans in the admission process, can be found in another case study. In 2013-14, the University of Texas, USA, introduced GRADE, an assistive ML recommendation system meant to speed up PhD candidates approval in the university.[26] Although the final admission decision was still to be made by humans, GRADE supplied the admission committee with a list of candidates who are near the decision boundary, based on historical admission data. In 2020 the computer science department of the University of Texas recognized that GRADE was exacerbating existing inequality in the department and phased out its use.[27,28] Like GRADE, more and more tools are being made available for institutions, promising to harness the power of ML to the process.[29,30] With such development, it is clear that further socio-technical discussion is required before widely adopting this technology.

In a related example from the educational system in the United Kingdom, ML was used to predict missing scores of exams. Due to the Covid-19 pandemic,[31] students in the UK missed their 2020 A-tests, which are normally used as admission exams into college. In their place, the government decided to plainly use an ML system, receiving minimal information about each student and the school they attend, to predict what the score would have been- and use that for admission. Although the average score around the country was higher than normal, many felt that this is an unaccepted approach. Indeed, the system demonstrated a bias in favor of students attending “rich” schools. [32,33] The incident strengthens the argument that a social discussion about ML-enhanced systems is mandatory.

4) Suitable socio-technical practices

In the previous chapters I defined and broke-down the academic admission process into four steps: Defining goals for the institution, locating suitable candidates, ranking them and finally, accepting the most suitable. I have presented some of the challenges that concern each step, with an emphasis on the third step, the ranking. Note that even though I have referred to the possible use of AI in the process, most of the mentioned problems apply regardless of whether or not AI is actually employed. To make the discussion more concrete and useful for those who seek to integrate AI in the admission process, this chapter reviews common socio-technical approaches and tries to adapt them accordingly.

Within the socio-technical realm, various suggestions could be made to bridge the gaps I discussed earlier. The umbrella term ‘Socio-Technical Systemic Design’ (STSD) encompasses a design approach which considers not only technical aspects of a system, but social and organizational aspects as well.[34] As could be deduced from the previous chapters, considering AI’s role in the academic admission process is a social and organizational matter. Therefore, STSD is a reasonable candidate approach to look at. Unlike in technical sciences, however, such socio-technical approaches for design are not rigorously proven to always guarantee a system’s improvement in a predicted way. Rather, they are brought as possible approaches one may adopt when encountering certain situations. It is still up to humans to continuously monitor the outcome of the process, as I shall demonstrate in this chapter and in the following one.

The main approach I will now look at is called Design for Values (DfV), meaning that the goals of the system at hand are explicit social values. Such values could be, in our case, fairness, equality, scientific progress etc. While DfV is a generic term, two common methodologies spell out the specificities one could follow. These two methodologies are called Value Sensitive Design (VSD)[35] and Participatory Design (PD)[36]. The rest of this chapter will elaborate on each of these methodologies in general, starting with a common requirement to both: a stakeholder analysis.

Stakeholder analysis

Both VSD and PD require knowing who the affected parties, called stakeholders, are. The analysis should spell out who they are (specific individuals, certain groups of people, general animal species, etc.) and their relation to the topic at hand. Such analysis is often made when developing a new or changing an existing system, policy or project.

Varvasovszky et al. propose a set of key questions and steps to be addressed in such analyses.[37] Here are the most relevant considerations for this paper:

1. What is the aim of the stakeholder analysis?
2. What is the context?
3. Identifying stakeholders.
4. For each stakeholder, evaluating their interest in the issue, their influence, their position/attitude and the impact that the issue has on them.

According to Varvasovszky, these tasks are best performed by involving the stakeholders themselves in the analysis.

Participatory Design

PD is the active involvement of the identified stakeholders in the design process. This methodology relies on the notion that, by cooperating with the affected parties, the final product would serve everyone better. The justification is that by asking the stakeholders for what they need directly, the developers address what is needed instead of what the developers would naturally resort to.

However, some valid concerns may arise about PD. Scariot et al.[38] and Nielsen[39] mention some pitfalls one should avoid while employing PD. First, it is resources consuming to find representative users for the process. Second, a participating stakeholder may fail to accept some inherent trade-offs about the final product, leading to dissatisfaction. Third, not allowing for user input to take form, as some designers "care about their baby too much". Fourth, allowing for too much stakeholder input may overcomplicate the system and add goals beyond the original ones. The fifth and last point is administrative, in which Scariot et al. suggest to not neglect the actual humans who contribute to the process, and to acknowledge their contribution as it is often taken for granted.

Value Sensitive Design

PD considers the stakeholders in general, be it any entity that has any sort of link to the product, and seeks their input. VSD turns to a different authority for defining the goals of the product: society. While the system may still have goals defined (either by a PD process or solely by the owner), VSD gives social and moral considerations a more central place in the design process.[40]

At the core of VSD is an iterative process[41], hinted at in previous chapters. Friedman[42] identifies three main parts to be repeated, called "investigations":

1. *Conceptual investigation:* This part includes a stakeholder analysis as well as values analysis. The latter may include specifying the desired values to be fulfilled by the final system, as well as values that get fulfilled as a "by-product".
2. *Empirical investigation:* For the many limitations of an isolated designer making conceptual assumptions on their own, an empirical phase is needed. In this part of the iteration, the developer must interact with the stakeholders and with the system to assess if the it fulfills its goals and to what extent.
3. *Technical investigation:* As different technologies are more suitable for certain purposes than others, the right kind of technical solution should be developed and evaluated. I mention here again the solutionism trap[16], as the most appropriate solution might not be technical at all.

Following these steps iteratively and continuously should yield, according to Friedman, a system which is loyal to the values defined in the first step. Do note that even though the original formulation calls for three "steps", these are simply actions which inform each other and can sometimes be parallel or even overlap.[43]

5) Evaluation of the different methodologies

In the previous chapter I have shown two STSD methodologies, which could be used to design an admission process which integrates AI. Some of the challenges in the admission process could perhaps be solved using some sort of AI. For other challenges it may not be best to mechanize their handling. It is not unreasonable to assume that such a complex social process would require human skills at some point, but it could also occasionally benefit from AI. To find the right balance, each challenge must be examined individually to determine how it should be dealt with. Only once we deem AI an appropriate solution to a certain challenge can we start to design the technical solution.

In one sentence, the task at hand is this: *Can AI enhance any of the four steps of admission process, which enhancements would those be, and how?* To answer that, I turn in this chapter to the methodologies detailed earlier.

Stakeholder analysis

Both PD and VSD require a stakeholder analysis. In order to dive deeper and demonstrate the application of these two methodologies, it would be best to have such analysis done on the admission process. For this paper I will progress with a hypothetical analysis. It is by no means comprehensive and could have other outcomes depending on the specific context. The analysis summary is shown in Table 1.

The aim of the stakeholder analysis in this case is to determine who are the parties that are influenced by the academic admission itself, and how they are linked to it.

The context for the analysis has several aspects. I am trying to set the framework for improving upon existing admission process, perhaps by integrating AI into it. To specify further, and hopefully get a more useful result, I refer to a very specific context: admission to a closed-capacity Bachelor degree program in a western country. Social context is therefore quite abstract and the scope of this paper restricts the discussion about it. The next step would be to determine who is interested in such a process.

To *identify the various stakeholders*, I resort to the question: who is affected by the process? Without yet detailing how it will be achieved, who are the parties that have any interest in the process? First and foremost, the applicants are influenced greatly from their ability to get feedback on their candidacy. Second, the admission staff who utilizes the final procedure after it is developed, is also impacted by changes made to the procedure. Third, the university administration is ultimately linked to everything happening within the institution. There are many other possible stakeholders that could be considered, but either have less interest or are less influenced by the process.

The final part of the stakeholder analysis is *discussing the approach each stakeholder has* towards the proposed change (AI integration).

Applicants are greatly affected emotionally as well as procedurally by the right decision being made or biases. Hopefully, they would benefit from AI integration into the process, but have very little influence on it eventually being so or not.

Also highly interested in the admission process is the *admission staff* itself. They not only have to execute the process according to plan, but are the first to absorb ricochets from frustrated applicants. They are likely to have a lot of influence over the process, although they are ultimately “only” supposed to design and execute the process in a way that serves the goals, as defined by institution’s administration.

The admission process is the first point of contact that a prospective student has with an academic institution. It may shape the student’s perception of the university and as such, this process and some aspects of it should be of interest to the *university’s administration*. They would want to provide prospective students with the best admission experience on one hand, but on the other they may not want their staff to be overburdened.

The analysis will be used to apply PD and VSD to the admission process, which I will now elaborate on.

Aim:	Detail who are the parties that are influenced by the academic admission process.				
Context:	(Increasing) Level of transparency in the admission process in western countries.				
Stakeholders	**Characteristics**				
	Involvement	*Interest*	*Influence/ Power*	*Position (supportive?)*	*Impact on stakeholder*
Students (accepted, rejected)	Applying, following up on their status. Receiving feedback.	Very high	Low	Supportive	High
Admission staff	Use the final process.	Very high	Very high	Supportive	High
University administration	Define goals to the process.	High	High	Supportive	Moderate

Table 1: A stakeholder analysis summary, for AI integration in the admission process.

Participatory Design

As the name suggests, PD encourages the developers to continuously involve different stakeholders in the process. In the case of admission process, this will be hard. Let us recall Scariot et al.'s[38] first issue with PD: high costs. Attending to the most affected stakeholder of the process, the applicants, and involving them in the development process, will be very difficult. Even if it was easy to find applicants who are willing to give feedback, it would still be hard to assert that it is representative. Involving the second most crucial stakeholder, admission staff, seems more feasible. It will be easier to have them attend throughout the entire process. However, this will only tell the developers one side of the story.

Due to the difficulty of gathering a continuous feedback from a central stakeholder, the applicants, I find PD to not be a suitable approach to making any changes to the admission process.

Value Sensitive Design

Choosing a different design approach does not obviate the need for input from applicants. Ignoring stakeholders in any socio-technical system design is not recommended by any methodology. Yet, engaging in VSD allows us to come back to different stakeholders (such as applicants) along the way in an iterative manner and receive feedback on the system's latest version each time. This iterative approach resembles the agile methodology,[44] a developing routine well-practiced among software developers.

For these two reasons (its relative simplicity compared to PD and the familiarity of developers with agile methodologies) I find VSD to be more suitable to apply to the admission process. The next chapter illustrates this application.

6) Application

Before deciding to apply a socio-technical methodology such as VSD to every step of the admission process, it would be wise to consider which steps of the admission process even have a technical aspect to them. To make sure that I use the right design tools and avoid the solutionism trap[16], I will only apply VSD to the third step: ranking candidates.

The VSD methodology calls for three analyses: conceptual, empirical and technical. I will now suggest a way to use these steps iteratively, together with the stakeholder analysis done earlier, in order to integrate AI in the candidates ranking step.

Conceptual investigation: stakeholder and values analysis

The hypothetical stakeholder analysis done in the previous chapter indicates three main stakeholders: the applicants, the admission staff and the institution's higher administration. In theory and from an organizational point of view, the administration should determine what the values that are to be fulfilled. For the sake of this paper, I will again propose some hypothetical values which could serve as guidelines. Recommendations for setting the institution's goals will be addressed further in chapter 7.

Applicants are likely to seek *equal opportunities*. This value was thoroughly neglected in both case studies shown earlier (University of Texas and the A-Level scores prediction in the UK). An applicant may be more willing to accept their rejection if they know that they had full control over the way the ranking system perceived them. In the UK case, for example, students' predicted grades were biased by to the school they came from and less by their past achievements, which naturally caused a major dissatisfaction. A student from a certain school had no control over their school's perception by the system, and as a result was not in control over their own representation in the eyes of the system.

The university's administration, from their side, would probably draw from the institution existing values: *scientific progress*, *human advancement*, etc.

Empirical investigation: stakeholder interaction

The developers should engage with the stakeholders themselves. This should be done either to survey what the values are and to assert that the most updated version of the system adheres to them. This step may include simulations, interviews and focus groups[45] with applicants, admission staff and administrators.

Technical investigation: which technology serves the values best?

I will focus on the value of *equal opportunities*, which was mentioned under conceptual investigation. This value directly relates to bias, to being aware of it and trying to avoid it. Defining bias and fairness is far from straightforward, even in the case study mentioned earlier of the 2020 A-Level grades in the UK.[33] As Wachter et al. point out, the algorithm did exactly what it was designed to do- predict likelihood of grades given particular student's scores and their school's past performance.[46] The bigger question is, is it fair? Once the matter was published, the grades were quickly amended to better match teacher's predictions. But this is was just a superficial fix to a symptom, not a cure for the underlying problem of the lacking definitions for fairness in that context.

Fixing all of society's inequality issues goes quite beyond the scope of this paper. Given these issues, one's best hope is to at least design a system in a way that is bias aware and continuously self-correcting. "System" in this case extends beyond the algorithm. It is also the people developing and using it. They too should be open to criticism and expect backlash. For example, once a certain ML ranking algorithm is chosen, it must be iteratively evaluated by the relevant stakeholders. As suggested by Davis et al., VSD is indeed an iterative, overlapping methodology.[43] Special attention should be given to common issues which often arise in these discussions about ML ranking of candidates: Bias, fairness and transparency. The technology used may change through time, but these principles should remain.

7) Discussion and reflection

In the suggested application, I created a hypothetical stakeholder analysis and proposed a way to use it with the VSD methodology to enhance the candidates ranking step of the admission process. This is merely an example of how the process could take place. The suggestions made focused on step 3 of the admission process, the ranking step. The first two steps, defining goals and locating candidates, could benefit from similar methodologies, perhaps leaning more towards social aspects.

The previous chapter illustrated very high-level suggestions and methodologies that could be followed to integrate AI into the ranking step. I shall now make some more concrete recommendations about the admission process as a whole and about education for ethical perspective in general, basing both on my own personal experience.

Further recommendations for the admission process

Previous recommendations dealt with the third step of the admission process, ranking candidates. However, the steps preceding it could benefit from a socio-technical methodology just as well. For the most part of the discussion, I showed how technical aspects needed to be enhanced by a social aspect. In the case of defining the institution’s goals, I believe that the opposite is true. Values and goals are considered to be a part of the social, maybe even philosophical realm. In my opinion, some technical perspective is required and could contribute to a better relation between the “soft” part of determining the institution’s goals and the later “harder” part of utilizing an algorithm to rank candidates by them. There are three important things to remember when defining goals:

1. First and foremost, the goals should correspond with the organization’s values. For example, many of the modern organizations in the western world would agree, at least on a declarative level, that they should disregard people’s gender while making organizational decisions concerning them. So, although it might seem trivial, I nevertheless mention it here for completeness that the goals of an academic institution should match its values. This is to either be done explicitly, or at least have the goals not contradict these values.
2. Of most relevance to this paper, while defining goals, one should give thought as to how these goals can be translated into formal requirements. This is to be done under the assumption that *“… the full meaning of social concepts such as fairness… cannot be resolved through mathematical formalisms.”*[16]
3. An institution should define its goals as clearly as possible, with as little room for personal interpretation as can be. This would decrease the chance of individuals abusing the goals as well as help the organization’s members to cooperate.

Defining and improving the institution’s goals should be an iterative process, as suggested by Van der Bijl-Brouwer[41]. I have expanded on this idea, employing methodologies from Value Sensitive Design[42], Participatory Design[16] and stakeholders analysis, in chapter 6.

My further recommendations about the admission process are brought here:

1. As illustrated by the UK A-level grades prediction case-study, personal data and achievements should play a more significant role in the admission decision than the contextual variables. In the mentioned case-study, only after major public backlash did the government prioritize personal

grades and teachers' recommendations over the general school average, finally allowing students to be evaluated individually.

2. Prediction currently has limitations. In my opinion, trying to predict the exact grade of a student is too ambitious for today's technology and society. It may be reasonable to determine which student is more likely to succeed than others in their academic studies, but the exact ranking should still be made human based.
3. In every process, even when machines are not involved and the decision is purely human-made, there is always an appeal process available for rejected applicants. They can, at the very least, receive an explanation for the rejection. In other cases, they may engage once again with the admission team and continue the discussion, sometimes reversing the admission decision. In my opinion, this functionality should be present especially in AI-enhanced admission processes. As explained in the introduction, this is not easy to implement and should be considered carefully.

Recommendations for ethical perspective in engineering education

The "classic" engineering studies focuses on "hard sciences". From mathematical courses to system design and project management, the social aspect seems to be pushed to the margins. The TU Delft in The Netherlands, for example, incorporates an ethics module in about one-in-six courses in the Bachelor program for computer scientists. This is meant to develop sensibility, analysis skills, creativity and more in the moral context.[47]

However, some studies argue that such educational efforts may not meet their original goal. These educational modules, not necessarily in the TU Delft, tend to focus on major cases that hit the news and/or only on the perspective of one party involved. They often lack social context and framing that would have given the students the full picture and contribute to their understanding of how complex real-life situations can be.[48,49]

In my opinion, a dedicated course about socio-technical design methodologies would contribute greatly to a more holistic approach in engineering education. Students would be encouraged to always consider moral and ethical implications of their designed systems.

Reflection about this work

This paper resembles a social study more than a technical paper. I believe that, as this is a part of a technical Bachelor program, there is merit to someone of such technical background writing about system design methodologies in a social context. This integration is the essence of the whole Socio-Technical approach.

Another critical point is to note that almost all studies and references used in this paper were made in either the United States or in the European Union. As such, they probably reflect the values and methodologies used in this part of the world and should never be considered for anything more. Further work must observe a wider variety of countries and institutions.

The objectiveness of this paper ends roughly halfway through. Breaking down the admission process in the beginning is merely a matter of definition, and the socio-technical methodologies review is a summary of common academic knowledge. From there, I have progressed to giving mostly my own ideas based on my experience and beliefs. That part should, again, be considered critically and further work could correct or expand on my conclusions.

8) References

1. Posted by Ehtesham Peerzade on December 11, 2019 at 2:30amView Blog. *Machine Learning Market is Rising Due to Rapid Increase in Unstructured Data*. *Data Science Central*.
2. Mitchel, T. *Machine Learning*. (1997).
3. Bau, D. *et al.* Understanding the role of individual units in a deep neural network. *Proc Natl Acad Sci USA* **117**, 30071 (2020).
4. AlGhamdi, A., Barsheed, A., AlMshjary, H. & AlGhamdi, H. A Machine Learning Approach for Graduate Admission Prediction. in *Proceedings of the 2020 2nd International Conference on Image, Video and Signal Processing* 155–158 (2020).
5. Srivastava, D. K., Sharma, A. K. & Shankar, V. G. Prediction of Admission in Master's Program in Foreign Countries Using Machine Learning Algorithms. *Data Engineering and Intelligent Computing: Proceedings of ICICC 2020* 239.
6. Acharya, A. & Sinha, D. Early prediction of students performance using machine learning techniques. *International Journal of Computer Applications* **107**, (2014).
7. Wachter, S., Mittelstadt, B. & Russell, C. Why Fairness Cannot Be Automated: Bridging the Gap Between EU Non-Discrimination Law and AI. *SSRN Electronic Journal* (2020) doi:10.2139/ssrn.3547922.
8. Aizenberg, E. & van den Hoven, J. Designing for human rights in AI. *Big Data & Society* 1–14 (2020) doi:10.1177/2053951720949566.
9. Cowgill, B. *et al.* Biased Programmers? Or Biased Data? A Field Experiment in Operationalizing AI Ethics. in *EC '20: Proceedings of the 21st ACM Conference on Economics and Computation* 679–681 (Association for Computing MachineryNew YorkNYUnited States, 2020). doi:10.1145/3391403.3399545.

10. Mittelstadt, B., Allo, P., Taddeo, M., Wachter, S. & Floridi, L. The ethics of algorithms: Mapping the debate. *Big Data & Society* 1–21 (2016) doi:10.1177/2053951716679679.

11. McCracken, M. Nine Forms of Bias in Admissions. *Kira Talent* https://blog.kiratalent.com/nine-forms-of-bias-in-admissions/ (2020).

12. Liem, C. C. S. *et al.* Psychology Meets Machine Learning: Interdisciplinary Perspectives on Algorithmic Job Candidate Screening. in *Explainable and Interpretable Models in Computer Vision and Machine Learning* (eds. Escalante, H. J. et al.) 197–253 (Springer International Publishing, 2018). doi:10.1007/978-3-319-98131-4_9.

13. Zehlike, M. *et al.* FA*IR: A Fair Top-k Ranking Algorithm. in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* 1569–1578 (Association for Computing Machinery, 2017). doi:10.1145/3132847.3132938.

14. Derous, E. & Ryan, A. M. When your resume is (not) turning you down: Modelling ethnic bias in resume screening. *Human Resource Management Journal* **29**, 113–130 (2018).

15. Trist, E. L. *The evolution of socio-technical systems*. vol. 2 (Ontario Quality of Working Life Centre Toronto, 1981).

16. D. Selbst, A., boyd, danah, A. Friedler, S., Venkatasubramanian, S. & Vertesi, J. Fairness and Abstraction in Sociotechnical Systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency* (2019) doi:10.1145/3287560.3287598.

17. Conrad, C. University Goals: An Operative Approach. *The Journal of Higher Education* **45**, 504–516 (1974).

18. TU Delft. TU Delft Strategic Framework 2018-2024. (2018).

19. Gross, E., Grambsch, P. V. & Schwab, J. J. University Goals and Academic Power. *The Annals of the American Academy of Political and Social Science* **384**, 175–176 (1969).

20. Greenwood, S., Perrin, A. & Duggan, M. Social media update 2016. *Pew Research Center* **11**, 1–18 (2016).

21. What We Look For - HARVARD. *Harvard College* https://college.harvard.edu/admissions/apply/what-we-look.

22. Fitzsimmons, W. R. Guidance Office: Answers From Harvard's Dean, Part 1. *The New York Times* https://thechoice.blogs.nytimes.com/2009/09/10/harvarddean-part1/ (2009).

23. Stanford, A. Our Selection Process - Stanford. https://admission.stanford.edu/apply/selection/#:~:text=The%20primary%20criterion%20for%20admission,is%20your%20high%20school%20transcript.

24. Nance, B. G. *et al.* What we look for- MIT. *MIT Admissions* https://mitadmissions.org/apply/process/what-we-look-for/ (2007).

25. Yale, A. What Yale Looks For. https://admissions.yale.edu/what-yale-looks-for#:~:text=The%20single%20most%20important%20document,school%20and%20have%20done%20well.

26. Waters, A. & Miikkulainen, R. GRADE: Machine Learning Support for Graduate Admissions. *AI Magazine* **35**, 64 (2014).

27. Austin, C. S. at U. *TXCS is deeply committed to addressing the lack of diversity in our field. We are aware of the potential to encode bias into ML-based systems like GRADE, which is why we have phased out our reliance on GRADE and are no longer using it as part of our graduate admissions process. Twitter* (Twitter, 2020).

28. Burke, L. The Death and Life of an Admissions Algorithm. *Inside Higher ED* (2020).

29. Ellucian & education, A. the A. E. I. by the transformative impact of. *Improving the enrollment process through machine learning*. *Ellucian*.

30. TargetX. *Request A Demo*. *TargetX*.

31. *Archived: WHO Timeline - COVID-19*. *Archived: WHO Timeline - COVID-19* (World Health Organization).

32. Piper, K. *The UK used a formula to predict students' scores for canceled exams. Guess who did well. Vox* (Vox, 2020).

33. Bedingfield, W. *Everything that went wrong with the botched A-Levels algorithm*. *WIRED UK*.

34. Baxter, G. & Sommerville, I. Socio-technical systems: From design methods to systems engineering. *Interacting with computers* **23**, 4–17 (2011).

35. Friedman, B. Value-sensitive design. *interactions* **3**, 16–23 (1996).

36. Schuler, D. & Namioka, A. *Participatory design: Principles and practices*. (CRC Press, 1993).

37. Varvasovszky, Z. & Brugha, R. A stakeholder analysis. *Health policy and planning* **15**, 338–345 (2000).

38. Scariot, C. A., Heemann, A. & Padovani, S. Understanding the collaborative-participatory design. *Work* **41**, 2701–2705 (2012).

39. Nielsen, J. Bridging the designer-user gap. *Useit. com* 17–03 (2008).

40. Van den Hoven, J., Vermaas, P. E. & Van de Poel, I. *Handbook of ethics, values, and technological design: Sources, theory, values and application domains*. (Springer, 2015).

41. Bijl-Brouwer, M. van der. Designing for Social Infrastructures in Complex Service Systems: A Human-Centered and Social Systems Perspective on Service Design. *She Ji: The Journal of Design, Economics, and Innovation* **3**, 183–197 (2017).

42. Friedman, B., Kahn, P. H. & Borning, A. Value sensitive design and information systems. *The handbook of information and computer ethics* 69–101 (2008).

43. Davis, J. & Nathan, L. P. Value sensitive design: Applications, adaptations, and critiques. *Handbook of ethics, values, and technological design: Sources, theory, values and application domains* 11–40 (2015).

44. Cohen, D., Lindvall, M. & Costa, P. An introduction to agile methods. *Adv. Comput.* **62**, 1–66 (2004).

45. Langford, J., Wilson, J. & Haines, H. Beyond focus groups? The use of group discussion methods in participatory design. in *Focus Groups* 159–174 (CRC Press, 2002).

46. Wachter, S., Mittelstadt, B. & Russell, C. Bias preservation in machine learning: the legality of fairness metrics under EU non-discrimination law. *West Virginia Law Review, Forthcoming* (2021).

47. Moral educational goals - TU Delft. https://www.tudelft.nl/ethics/ethics/the-delft-approach/educational-goals.

48. Harris Jr, C. E., Davis, M., Pritchard, M. S. & Rabins, M. J. Engineering ethics: what? why? how? and when? *Journal of Engineering Education* **85**, 93–96 (1996).

49. Newberry, B. The dilemma of ethics in engineering education. *Science and Engineering Ethics* **10**, 343–351 (2004).