The influence of algorithm aversion and anthropomorphic agent design on the acceptance of AI-based job recommendations

Completed Research Paper

Jessica Ochmann[1]
jessica.ochmann@fau.de
Leonard Michels[2]
leonard.michels@fau.de
Sandra Zilker[3]
sandra.zilker@fau.de
Verena Tiefenbeck[2]
verena.tiefenbeck@fau.de
Sven Laumer[1]
sven.laumer@fau.de

Friedrich-Alexander-University Erlangen-Nuremberg, Germany
[1]Schöller Endowed Chair for Information Systems,
Fürther Str. 248, 90429 Nueremberg
[2]Digital Transformation Group,
Lange Gasse 20, 90403 Nuremberg
[3]Chair of Digital Industrial Service Systems,
Fürther Str. 248, 90429 Nuremberg

Abstract

Artificial intelligence (AI) offers promising tools to support the job-seeking process by providing automatic and user-centered job recommendations. However, job seekers often hesitate to trust AI-based recommendations in this context given the far-reaching consequences of the importance of the decision for a job on their future career and life. This hesitation is largely driven by a lack of explainability, as underlying algorithms are complex and not clear to the user. Prior research suggests that anthropomorphization (i.e., the attribution of human traits) can increase the acceptance of technology. Therefore, we adapted this concept for AI-based recommender systems and conducted a survey-based study with 120 participants. We find that that using an anthropomorphic design in a recommender system for open positions increases job seekers' acceptance of the underlying system. However, algorithm aversion rises if detailed information on the algorithmic origin is being disclosed.

Keywords: *Algorithm aversion, Anthropomorphism, AI-based recommendations, Human Resource Management*

Introduction

On average, we spend approximately a quarter of our lives working (Pryce-Jones 2010) and our work life significantly contributes to our level of well-being (Bowling et al. 2010). Consequently, the decision which jobs to apply to is of major importance. To find a position in line with their expectations, job seekers typically screen plenty of job proposals and apply for the most appropriate ones, putting much effort into optimizing their applications (Berg et al. 2010); this process, however, can be stressful and time-consuming (Wanberg et al. 2010). Recent developments in artificial intelligence (AI) bring promising avenues for overcoming these issues by providing automatic, rich, and user-centered recommendations, holding the potential to considerably improve the way individuals search and apply for jobs.

This approach is increasingly adopted in the human resources (HR) context, which has fostered the development of AI-based job recommender systems (Duan et al. 2019). These systems pre-select job alternatives based on job seekers' preferences, personal data, and data from previous job-seekers with comparable profiles, by using machine learning (ML) algorithms that predict the most suitable vacancies (Castelvecchi 2016). The resulting AI-based recommendations support job seekers in finding job proposals that best fit their individual preferences and qualifications (Malinowski et al. 2006). Due to these promising results, these systems are very likely to greatly impact the future of recruiting both from an organizational and from an individual perspective. However, ultimately the success of job recommender systems hinges on a large number of job seekers using them. Consequently, a better understanding of the factors determining the acceptance of the recommendations provided by these systems is required. In fact, many individuals still hesitate to rely on them (Laumer et al. 2018). While research in this field in general, and the recruitment domain in particular, is still scarce (van Esch et al. 2019), the existing body of literature on the acceptance of recommendations in the general business-to-consumer (B2C) context offers a valuable point of departure to build upon.

To uncover the factors that govern the acceptance of recommendations, prior research has typically used frameworks that incorporate cognitive aspects (e.g., Hu and Pu 2009) or relational constructs such as trust (e.g., Komiak and Benbasat 2006). However, these frameworks and this approach in general may not necessarily apply to high-stake contexts like job-seeking. The main difference between recommendations in the general consumer context and the job-seeking domain is that the decision to rely on a recommendation for a new job will have much more important implications on one's life, compared to simple, inconsequential decisions such as choosing which movie to watch or which song to listen to. Therefore, it is very plausible that algorithmic aversion (Dietvorst et al. 2015; Logg et al. 2019) arises when individuals encounter AI-based recommendations in high-stake decision contexts.

In general, algorithm aversion describes the phenomenon that individuals are often reluctant to accept and to rely on results computed by statistical algorithms and rather trust human forecasts, even though evidence-based algorithms are more accurate in predicting appropriate alternatives compared to human reasoning (Dietvorst et al. 2018). Consequently, many individuals are hesitant or entirely refuse to rely on recommendations that are apparently based on algorithmic prediction (Burton et al. 2020; Castelo et al. 2019; Dietvorst et al. 2015). This tendency might be especially prevalent in high-stake decisions supported by algorithm-based recommendation systems. As humans often assume that they have superior reasoning compared to algorithms (Dietvorst et al. 2015, 2018), relying on an algorithm's recommendation is perceived as a more risky decision than relying on one's own reasoning. Prior research has shown that when the stakes of a decision rise, humans tend to become risk-averse (Fehr-Duda et al. 2010), which would manifest in algorithm aversion in high stake decisions. Recent research, indeed, suggests a two-sided character of this phenomenon, showing that algorithm aversion is not omnipresent and can be reduced by giving users more control and allowing them to modify the algorithm (Dietvorst et al. 2018). Further, in situations that require ample background knowledge (e.g., prediction of business or geopolitical events), users even display a certain level of algorithm appreciation (Logg et al. 2019).

To overcome the potential issue of algorithm aversion, it might be beneficial for AI-based recommender systems not to reveal the algorithmic origin of their recommendations. Users might be overwhelmed by the complex information or mistrust it and consequently decide to rather rely on their judgment of which the underlying processes appear clearer to them. In addition to this potential measure to avoid effects decreasing the acceptance of AI-based job recommendations, it seems beneficial to investigate the effects of measures that have been found to increase the acceptance of AI-based recommendations in this high stake decision context. Prior research emphasizes that users are more likely to accept the choice of a recommender system when it is presented as a human-like agent (Qiu and Benbasat 2009). One unobtrusive, easy-to-implement measure to increase the human-likeness of a recommendation agent could be the use of anthropomorphic (human-like) design features for the presentation of the recommendation (Epley et al. 2007; Pfeuffer et al. 2019; Qiu and Benbasat 2009; Wang et al. 2016). Hence, this study focuses on the investigation of why and when job seekers will adopt AI-based recommender systems. Our aim is to investigate whether algorithm disclosure and anthropomorphism influence the acceptance of AI-based recommendations in the job-seeking context. We thus contribute to the literature on algorithm acceptance research (Dietvorst et al. 2018; Logg et al. 2019) by empirically investigating the effect of algorithm disclosure and anthropomorphism on the acceptance of AI-based job recommendations, which is a more serious high stake context than the ones examined in prior research. Thus, we answer the following research question:

> ***RQ:*** *How does disclosing detailed information on the algorithmic origin of an AI-based job recommendation and the use of anthropomorphic design features to communicate an AI-based job recommendation affect its acceptance by users?*

The paper is organized as follows. First, we briefly discuss relevant literature on recommender systems in human resource management, algorithm aversion regarding AI-based recommendations in particular, the use of anthropomorphism in recommendation systems, and derive our hypotheses. Next, we motivate our choice of the scenario-based technique used in the present study and describe the methodology we used to test our hypotheses. Finally, we report the results of our empirical study and conclude with a general discussion of the findings.

Related work and development of hypotheses

Recommender systems in human resource management

Recommender systems were first introduced by Resnick et al. (1997) and describe information systems that analyze user data to produce personalized recommendations that match user preferences. Their objective is to reduce a user's potential information overload by sorting and filtering alternatives in terms of relevance and user fit. Besides the benefits provided for the user, effective recommender systems also help organizations that offer these systems to increase consumer loyalty and sales and to differentiate themselves from competitors (Adomavicius et al. 2019; Gomez-Uribe and Hunt 2015; Sharma et al. 2015). Over the last decade, the application of recommender systems has noticeably increased in a variety of domains, such as e-commerce, media, and human resources (Lu et al. 2015; Malinowski et al. 2006). While recommender systems in e-commerce and media predominantly aim to reduce consumers' efforts necessary to find relevant products or services, in the recruiting context, two different types of recommender systems are discussed that address either the organization or the job seeker. First, in the organizational context, CV recommender systems that are used by recruiters to match a specific vacancy with the most appropriate candidate; second, job recommender systems for job seekers that match their job preferences with suitable vacancies (Malinowski et al. 2006).

Given the omnipresence of recommender systems, their growing importance in individual decision-making processes, and their large economic potential (Bodapati 2008), it is crucial for academia and practitioners alike to understand the factors that influence the acceptance of these systems. Therefore, scholarly research put increasing effort in the investigation of various theories and models to explain the acceptance of recommender systems and their results, with a strong focus on the domain of product recommendations in e-commerce (Adomavicius et al. 2019; Komiak and Benbasat 2006; Moussawi et al. 2020).

It remains unclear, however, if and how these results can be adapted to the job-seeking context that is characterized by high personal stakes of the decision, as job satisfaction highly influences life satisfaction (Bowling et al. 2010). In addition, prior research highlighted especially the acceptance of conventional job recommender systems that rely on collaborative filtering and content- or knowledge-based techniques (for a review see Lu et al. 2015; Park et al. 2012). Currently, the diffusion of and advances in the research on artificial intelligence provide additional opportunities for recommender systems to make more precise and user-centered recommendations (Dietvorst et al. 2018). Algorithms gain in prediction quality, and the resulting AI-based recommendations have the ability to assist users in the preselection of alternatives in a more sophisticated manner as they are able to discover intricate structures in large data sets (LeCun et al. 2015). Thus, AI-based recommender systems have the potential to fundamentally change future job search. Therefore, our aim is to unveil factors that influence the acceptance of AI-based recommendations. In line with prior research (Promberger and Baron 2006), we introduce the acceptance of AI-based recommendations as our dependent variable of interest. In the following, we discuss the concepts of *algorithm aversion* and *anthropomorphism* as potentially influential factors regarding the acceptance of AI-based job recommendations.

Algorithm aversion

Algorithms are defined as computer-implementable instructions to perform a determined task and often outperform human experts in various tasks (LeCun et al. 2015). Multiple scholars have theorized an *aversion towards algorithms* that give users automated advice regarding a certain task (Castelo et al. 2019; Dana and Thomas 2006; Dietvorst et al. 2018) although pioneering research from the 1950s illustrated that even basic statistical algorithms such as linear regression outperform human experts on medical diagnosing tasks (Dawes et al. 1989; Meehl 1954). Since then, the fast progress in the field of artificial intelligence has enabled algorithms to learn from the past, understand and create natural language, and even reflect human emotions (Castelo et al. 2019), further increasing their potential superiority compared to human reasoning.

The increasing presence of algorithms, however, confronts individuals more frequently with the choice of whether they should rely on human experts or on algorithms. The dominant theme in this broad academic research area is that individuals prefer human advice over algorithms (Dietvorst et al. 2015). The underlying reasoning of this so-called *algorithm aversion* is manifold and can be ascribed to the desire for a perfect prediction and the mistaken belief that humans are more capable of perfection (Einhorn 1986), ethical concerns (Dawes 1979; Eastwood et al. 2012), and the zero error tolerance for algorithms (Dietvorst et al. 2015). Moreover, the lack of perceived control over the forecast inhibits the acceptance of algorithms (Dietvorst et al. 2018). Scholars have further argued that individuals' mistrust towards machines results in rejection of algorithm advice (for a review see Castelo et al. 2019). For example, individuals assume that an algorithm is unable to take one's unique circumstances fully into account and are averse regarding automated medical care (Longoni et al. 2019). In the field of recruitment, Diab et al. (2011) found that participants expect human recruiters to be more useful, professional, fair, and flexible than algorithms that are programmed to select employees.

In contrast, recent scholarly work reflects that for numerical tasks with an objectively correct answer, individuals actually prefer advice from algorithms to advice from human beings. This phenomenon is subsumed under the term *algorithm appreciation* (Logg et al. 2019). In addition, algorithm familiarity for a certain task increases trust and acceptance of algorithms. For example, individuals who are familiar with product or movie recommendations on according platforms tend to rely on the advice of these algorithms (Castelo et al. 2019).

These controversial findings highlight the need for further academic research to unveil reliable factors that predict the acceptance of algorithms for different types of tasks and contexts (Castelo et al. 2019). The systematic exploration of why and when individuals accept algorithms further helps to build an understanding under which circumstances job seekers rely on AI-based recommendations. As prior research suggests that disclosing information on the algorithmic origin of a recommendation might lead to reluctance regarding its acceptance (Burton et al. 2020; Castelo et al. 2019), our study seeks to induce algorithm aversion by varying the amount of information provided on the algorithmic origin of an AI-based job recommendation. We call this manipulation algorithm disclosure to address the user's potential algorithm aversion when being directly confronted with the advice of an algorithm. In line with prior research, we assume that

Hypothesis 1: *Disclosing the information on the algorithmic origin of an AI-based job recommendation leads to a lower acceptance rate of this recommendation.*

While algorithm aversion is a concept that can explain why users refrain from accepting AI-based job recommendations, it does not provide a potential lever to actively increase the acceptance rate of such recommendations. As prior research has shown, individuals rather trust humans than algorithms when it comes to recommended decisions (Dietvorst et al. 2018; Qiu and Benbasat 2009). One promising approach to increase the acceptance of AI-based job recommendation systems could, therefore be, to increase the human-likeness of the system. This approach has been discussed by prior research under the term of anthropomorphism.

Anthropomorphism

The concept of anthropomorphism refers to the process of attributing human characteristics, traits, or features to non-human agents, in order to reduce uncertainty and increase comprehension in situations when knowledge about the mechanisms underlying the behavior and the intentions of the non-human agent is scarce (Epley et al. 2007; Pfeuffer et al. 2019). By anthropomorphizing the non-human agent, users make inferences about themselves and other humans to predict its future behavior or make sense of its past behavior. If the evaluation of the anthropomorphized non-human agent is positive, it will be associated with multiple other positive characteristics such as trustworthiness, reliability, or competence (e.g., Aggarwal and McGill 2007; Benlian et al. 2019; Mourey et al. 2017; Qiu and Benbasat 2009; Wang et al. 2016). Prior research suggests that the positive effect of anthropomorphism on acceptance of a non-human agent is driven by an increased social presence, referring to the capacity of a technology to convey relational information and the extent to which it builds up a psychological connection with the user (Cyr et al. 2007, 2009; Qiu and Benbasat 2009; Schultze and Brooks 2019; Short et al. 1976). Thus, manipulating the extent to which a job recommendation system incorporates anthropomorphic design features seems to be a valuable and promising approach to increase its acceptance and the adoption of its recommendations.

The more human-like a non-human agent appears with respect to its visual, auditive, or mental characteristics, the more likely it will be anthropomorphized (Pfeuffer et al. 2019). The positive effects of this anthropomorphization might, however, disappear once the non-human agent becomes too lifelike, such that it raises a feeling of eeriness in the user that leads to revulsion - an effect that is known as the *uncanny valley* (Mori et al. 2012). A design manipulation that is easy to implement and efficient is the implementation of human images with facial features (Cyr et al. 2009; Gong 2008; Pak et al. 2012; Riegelsberger et al. 2003; Wang et al. 2016). Prior research has shown, for example, that the use of human images with facial features leads to more positive evaluations of websites and recommendation agents (Cyr et al. 2009; Pak et al. 2012; Steinbrück et al. 2002; Wang et al. 2016). Further, Qiu and Benbasat (2009) show that increasing the social presence of a product recommendation agent leads to an increase in the user's intention to use it as a decision aid. Pak et al. (2012) also report an increased adoption of a decision aid in a medical context when it was equipped with anthropomorphic characteristics. For personal intelligent agents (PIA), Moussawi et al. (2020) identified perceived anthropomorphism as an antecedent of PIA adoption. Gruber et al. (2018) and Gruber (2018), however, investigated the effect of anthropomorphism on the acceptance of navigation decision aids and found no significant effects. These results suggest that anthropomorphic design features can lead to increased recommendation acceptance, but also that this effect is not unequivocal. Further, research has not yet investigated the effects of anthropomorphic design features on the acceptance of recommendations in areas where high personal stakes are involved and one typically does not rely on automated recommendation systems, such as the job-seeking domain. To evaluate whether using anthropomorphic design features can increase the acceptance of AI-based job recommendations, we will manipulate whether the recommendation is communicated to the users using a human image or an artificial non-human image. In addition, in the anthropomorphic condition, we will refer to the recommendation agent in the first person, giving it a name and a gender, as prior research has shown that this further increases the degree to which users think of a non-human agent in human terms (Aggarwal and McGill 2007; Mourey et al. 2017). Based on the results of prior research, we assume that

Hypothesis 2: *Communicating the results of an AI-based job recommendation system using anthropomorphic design features leads to a higher acceptance rate of this recommendation.*

As prior research has shown, trust in the recommendation agent can have substantial effects on adoption behavior (Komiak and Benbasat 2008, 2006; Qiu and Benbasat 2009; Wang et al. 2016). These effects, however, do not necessarily persist when manipulating anthropomorphic design features or the degree to which an AI-based recommendation system is perceived as an elaborate algorithm (i.e., intelligent), as a recent study has shown (Moussawi et al. 2020). To account for potential trust effects on the effectivity of the above-mentioned interventions on the acceptance of AI-based recommendations, we control for general trust in artificial intelligence regarding job-related decisions in our analyses.

Research method

Experimental Design

We implemented the study as a two-factorial ((anthropomorphic: yes/no) x (AI process disclosure: yes/no)) between-subject design. More precisely, we manipulated A) whether the artificial intelligence was presented in an anthropomorphic way and B) whether the AI's underlying processes were disclosed to operationalize algorithm aversion. In the anthropomorphic condition (AN), we referred to the artificial intelligence in the third person and gave it a human name (i.e., "Emily"). Further, the AI-based recommendation was communicated to the participants by a picture of a woman (see Figure 1). We conducted a pre-test to ensure that the picture we used in the anthropomorphic condition did not appear negative on relevant dimensions (Wang et al. 2016). In this pre-test with N = 48 participants, we evaluated whether the person in the picture appeared professional, authoritarian, like an expert, trustworthy, dependable, reliable or like an HR expert using a 7-point Likert scale ranging from "Strongly disagree" to "Strongly agree"(Wang et al. 2016). The sample of participants of this pre-test was not from the sample pool of subjects as the participants in the final study and was recruited using different online sampling methods (e.g., social media groups, professional networks). Along all dimensions, the picture was assessed as significantly positive (i.e., compared to the neutral value of 4). In the non-anthropomorphic conditions ($\overline{\text{AN}}$), the artificial intelligence was not referred to in the third person, and its recommendation was communicated by a mechanical, abstract picture of gears (see Figure 1). In the two conditions with a high degree of algorithm disclosure (AD), participants were informed that the AI used algorithms, equations, and a comprehensive database to identify the most suitable position. In the two conditions with a low degree of algorithm disclosure ($\overline{\text{AD}}$), no such information was provided. Participants were assigned to one of the four conditions (i.e, AN_AD, $\overline{\text{AN}}$_AD, AN_$\overline{\text{AD}}$, $\overline{\text{AN}}$_$\overline{\text{AD}}$), using block randomization. Due to the involvement of human subjects, the authors seeked for and were granted approval for the study by the person at the university department overseeing the good conduct and ethical aspects of empirical research.

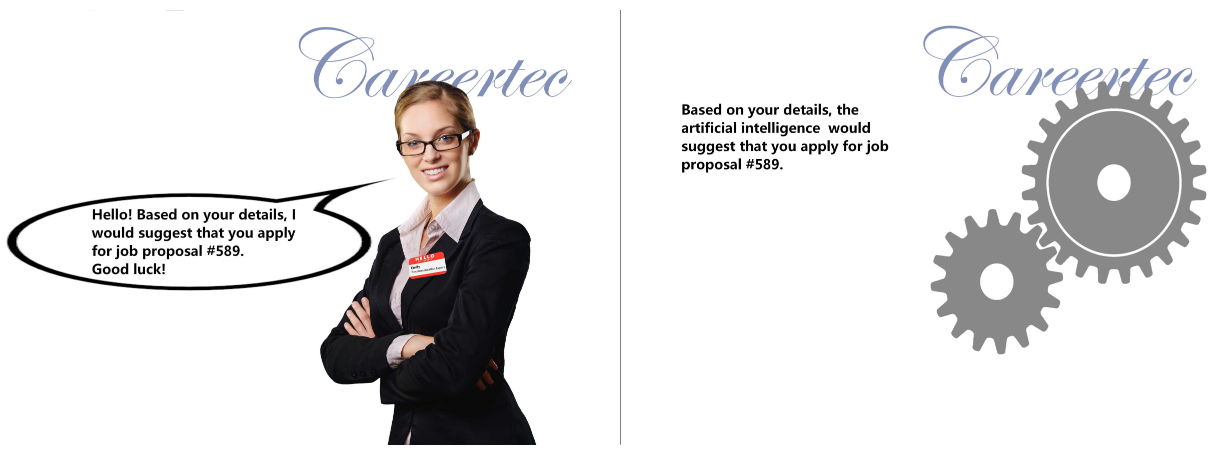

Figure 1. Pictures used to present the AI-based recommendation in the anthropomorphic condition (on the left), and non-anthropomorphic condition (on the right).

User study

To test our hypotheses, we applied a scenario-based vignette technique (Finch 1987) and conducted an on-line, survey-based user study. Further information on the recruitment procedure and study participants will be provided below. To protect the study's participants, they were informed about the content of the study, about its scientific background, and the measures taken to protect their privacy (i.e., anonymization) prior to starting the study. In addition, the authors' contact information was provided, allowing the participants to could contact them if they had any concerns related to the study or their personal data. By participating in the study, all participants agreed with this. The welcome page of the survey informed participants that the survey related to the job application context. Next, they were asked to put themselves in the situation that they were at the beginning of their thirties, unsatisfied with their current job position and, therefore, looking for a new full-time job for young professionals. This negative framing of their current job situation was used to increase the stakes associated with finding a new one, hence emphasizing the importance of the job searching process. To build up a certain level of initial trust (Moussawi et al. 2020), it was revealed to them that they, by chance, had learned about a career platform that has a very good reputation, successfully placed many job seekers, and was recently equipped with an AI-based recommendation system. The specific wording was used to induce trust in the platform and thus a baseline level of general trust in the AI based on the reputation of the platform. No direct trust-inducing measure for the AI was used to avoid interference with potential algorithm aversion. The recommendation system was described according to the experimental condition, and the participants were informed that they had only trial access to the platform, which meant that they could only apply to one of the suggested open positions. Participants were then asked for their first name, gender, highest degree or completed level of education, the kind of company they would like to work for, the department they would like to work in, and the city in which they were looking for a job. After these questions, they were shown an example of how the job proposals would be presented to them. Afterwards, a loading screen appeared for five seconds to imply that the AI-based recommender system was searching for suitable job proposals. On the next page, participants were presented with the results of the career platform and the job recommender system in the form of a text describing the results and procedure, according to the experimental condition, along with a picture stating the AI's recommendation (see Figure 1). Below this information, four different job proposals were provided to them, one with a blue frame and a yellow badge, representing the recommended job. The job proposal that was recommended to the participants was fixed across conditions (see Figure 2). The recommendations matched the department they wanted to work in. Participants were asked to consider the options for at least 60 seconds and were not able to proceed with the survey before that time had elapsed. On average, participants spent 107 seconds on the decision.

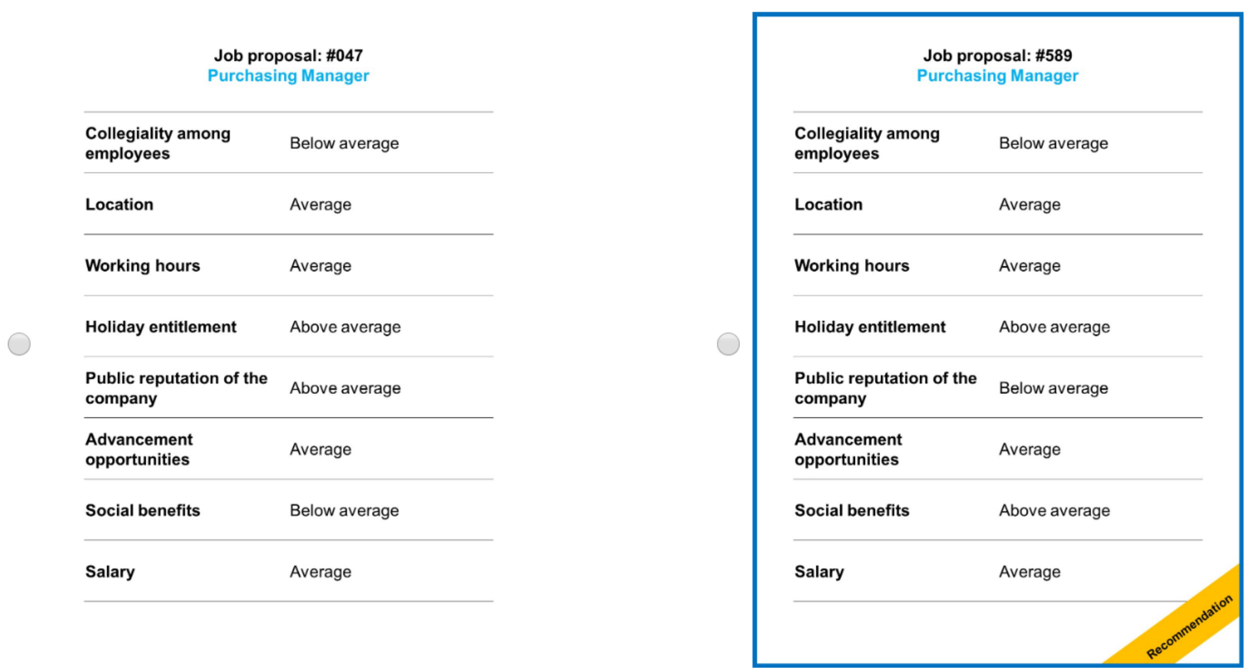

Figure 2. Two out of four job proposals presented to the participants in the experimental task, with one recommended proposal (on the right).

The participants' task consisted in choosing one of the four open positions they would like to apply for. The open positions were rated along eight dimensions as either average, above average or below average compared to jobs in other companies of the same industry. We opted for this approach, as prior research has shown that a choice set of four options described on eight dimensions leads to a suitable level of decision difficulty in assessing factors that may influence individual decision-making (Dijksterhuis et al. 2006). To determine which eight dimensions to use for describing the job proposals, we had conducted a second pre-test with 37 participants. We asked the participants of that pre-test to rank a set of 12 job dimensions by how important they perceived them in determining whether they would apply for a job in a comparable scenario to the one we used in our study; the mean ranks of the different dimensions are presented in Table 1. The participants of the second pre-test were recruited using online sampling methods (e.g., social media groups, professional networks); none of them had participated in the first pre-test or was part of the sample that participated in the final study. To avoid that one of the job proposals would be considered as the obvious choice by the participants, we selected the four most important dimensions according to our pre-test and ranked all job proposals as 'average' regarding them. In a second step, we selected four additional job dimensions that had been ranked as comparably important in the pre-test. Therefore, we selected the dimensions ranked from 6 to 9 in the pre-test. Regarding these dimensions, the job proposals differed such that for every job proposal, two dimensions were ranked as 'above average' and two as 'below average', while controlling that all four job proposals were ranked differently on these dimensions. After stating their decision, the participants were asked multiple questions regarding their decision and their impression of the recommendation system. Participants spent on average 448 seconds to fill out the post-task questionnaire.

#	Dimension	Mean rank (SD)
1	Salary	2.43 (1.59)
2	Advancement opportunities	4.30 (2.70)
3	Working hours	4.30 (2.20)
4	Location	4.68 (3.00)
5	Further education opportunities	5.89 (3.20)
6	Collegiality among employees	6.76 (2.85)
7	Public reputation of the company	6.89 (3.45)
8	Holiday entitlement	7.27 (2.59)
9	Social benefits	7.65 (2.87)

Table 1. Pre-test results (*N* = 37): Perceived importance of job dimensions in the application decision, mean rank and standard deviation (rank 1 = most important).

Scales and measurement variables

To receive a more detailed picture of the effects of anthropomorphic design features and algorithm aversion on the acceptance of AI-based job recommendations, we used multiple literature-based scales. The perceived human-likeness of the recommendation agent was assessed using the anthropomorphism scale by Bartneck et al. (2009); two items of the original scale were not included in the survey, as they referred to human-robot interaction. This resulted in three final items on which the participants' responses were assessed using 7-step semantic differentials. In addition, we used the technophobia scale of Sinkovics et al. (2002), consisting of four items, assessed on a 7-point Likert scale, to control for possible randomization artifacts regarding the prevalence of technophobia in the experimental groups. We adopted the scale to our scenario, as recent research suggests that technophobia can have a great impact on the adoption of new technology (Khasawneh 2018). We further evaluated the general trust participants had in artificial intelligence regarding job-related decisions by asking them to which extent they trusted the opinion of artificial intelligence when it comes to decisions about their professional future (using a 7-point Likert scale ranging from "not at all" to "very much").

Participants

The data was collected using the online participant recruitment service Prolific (Palan and Schitter 2018; Peer et al. 2017) with an English-speaking sample predominantly from the UK and the USA. We recruited 128 participants. Due to incomplete data, implausible overall duration of the experiment, and inconsistent answers, the data of 7 participants had to be excluded. A demographic summary of the final participants is provided in Table 2.

Sample	**AN_AD (N = 30)**	**$\overline{\text{AN}}$_AD(N = 30)**	**AN_$\overline{\text{AD}}$ (N = 30)**	**$\overline{\text{AN}}$_$\overline{\text{AD}}$ (N = 31)**
Gender				
Men	9 (30%)	13 (43.3%)	5 (16.7%)	14 (45.2%)
Women	21 (70%)	17 (56.7%)	25 (83.3%)	17 (54.8%)
Mean age (years)	30.1 (*SD* = 7.17)	32.3 (*SD* = 5.92)	29.5 (*SD* = 6.53)	31.8 (*SD* = 6.95)
Education level				
Primary education	2 (6.7%)	3 (10%)	5 (16.7%)	2 (6.5%)
Secondary education	7 (23.3%)	6 (20%)	9 (30%)	11 (35.6%)
Vocational training	1 (3.3%)	1 (3.3%)	0 (0%)	0 (0%)
University, undergraduate	20 (66.7%)	13 (43.3%)	10 (33.3%)	9 (29%)
University, postgraduate	0 (0%)	5 (16.7%)	4 (13.3%)	9 (29%)
Employment status				
Employed	25 (83.3%)	23 (76.7%)	20 (66.7%)	20 (64.5%)
Unemployed	1 (3.3%)	2 (6.7%)	3 (10%)	2 (6.5%)
Self-employed	0 (0%)	2 (6.7%)	3 (10%)	3 (9.7%)
Homemaker	2 (6.7%)	2 (6.7%)	1 (3.3%)	3 (9.7%)
Student	2 (6.7%)	1 (3.3%)	3 (10%)	3 (9.7%)

Table 2. Descriptive results of key socio-demographic data of the study sample.

Results

Randomization check

To ensure that the random assignment in the survey led to a uniform distribution of demographic criteria in the treatment groups, randomization checks were conducted for the four demographic indicators reported in Table 2. While the difference in the gender distribution between the four groups was marginally significant ($\chi^2(3) = 7.13$, $p = .068$), a subsequent post-hoc test adjusting p-values by the Benjamini–Hochberg procedure for multiple comparisons (Benjamini and Hochberg 1995) revealed no significant differences between the subgroups.
The treatment groups did not differ significantly regarding their mean age ($F(3, 117) = 1.21$, $p = .309$). While Fisher's exact test revealed significant differences in the distribution of the education level ($p = .020$), a descriptive interpretation of the results does not indicate any tendencies in the groups, however. The employment status distribution did not differ significantly between the groups ($p = .839$). Regarding the prevalence of technophobia, the four experimental groups did not differ ($M = 4.04$, $SD = 1.35$; $F(3, 117) = 0.91$, $p = .440$) on the technophobia scale by Sinkovics et al. (2002). For general trust in artificial intelligence regarding job-related decisions, we did also not find significant differences between the groups ($M = 4.06$, $SD = 1.33$; $F(3, 117) = 0.53$, $p = .660$).

Manipulation check

To assess whether the presentation of the recommendation system was perceived as anthropomorphic, we used the anthropomorphism scale developed by Bartneck et al. (2009) with a seven-point semantic differential. To adjust the scale for non-robot interaction and to keep the completion time of the survey to a reasonable limit, we had included only a subscale of the scale in the survey, excluding two items (i.e., Unconscious - Conscious; Moving rigidly - Moving elegantly). The scale showed sufficient internal consistency (Cronbach's $\alpha = .87$). Contrary to our assumptions, the AN groups did not perceive the recommendation system as more anthropomorphic than the $\overline{AN}$ groups (($t(119) = -0.24$, $p = .810$).

Hypothesis testing

To test our hypotheses, we used logistic models to assess whether the acceptance of the recommendation (dependent variable; dichotomized such that a value of 0 means that the recommendation was not chosen and 1 that the recommendation was chosen) was influenced by algorithmic disclosure and anthropomorphic design features (independent variables). This method was chosen as it allows to independently examine the effects of multiple predictors on a discrete outcome variable (Hosmer et al. 2013). In a first step, we included the algorithm disclosure variable to assess whether providing information on the process of how the AI determined the recommendation had an effect on the participants' decision to accept the recommendation. The variable was binary coded such that a value of 0 represented no algorithm disclosure and 1 algorithm disclosure. The logistic regression revealed a significant effect of the algorithm disclosure manipulation on the acceptance of the recommendation ($R^2_{Nagelkerke}$ = 0.06; $\chi^2(1) = 5.22$, p = .022), see Table 3. Disclosing information on the algorithmic disclosure of the recommendation significantly reduced the likelihood that the recommendation was chosen. This result is in line with our first hypothesis.

To investigate the effect of anthropomorphic design (independent variable) on the recommendation acceptance (dependent variable), we included whether participants received the recommendation in an anthropomorphic design in the logistic model as a dummy variable in a second model. The variable took a value of 0 for no anthropomorphic design and 1 for anthropomorphic design. The logistic regression revealed a marginally significant effect of the anthropomorphism manipulation on the acceptance of the recommendation ($R^2_{Nagelkerke}$ = 0.10; $\chi^2(2) = 8.98$, p = .011), see Table 3. The second model was marginally better in predicting the acceptance of the recommendation than the first model ($\chi^2(1) = 3.77$, p = .052). Participants exposed to the recommendation in an anthropomorphic design were more likely to follow the recommendation. This result has to be interpreted with caution, however, as it was only marginally significant. Our second hypothesis is therefore only partially supported.

In our third model, we further added the trust participants had in AI-based recommendations when it comes to decisions about their professional future as a moderator of the effect of algorithm disclosure and anthropomorphic design on the acceptance of the recommendation. This was based on the results of prior research (Moussawi et al. 2020). The logistic regression revealed no significant main effect of general trust in artificial intelligence and marginally significant interaction effects with algorithm disclosure and anthropomorphic design on the acceptance of the recommendation ($R^2_{Nagelkerke}$ = 0.23; $\chi^2(5) = 23.04$, $p < .001$), see Table 3. When adding general trust and its interaction with the effects of algorithm disclosure and anthropomorphic design features to the model, however, the effect of anthropomorphic design features on the recommendation acceptance became insignificant. The third model was significantly better in predicting the acceptance of the recommendation than the second model ($\chi^2(3) = 14.05$, p = .003).

	Coefficients	**Estimate (SE)**	***z*-value**	***p*-value**	**Odds Ratio [95%-CI]**
Model 1	Intercept	0.48 (0.27)	1.79	.073	
	Algorithm disclosure	-0.84 (0.37)	-2.26	.024*	0.43 [0.21; 0.89]
Model 2	Intercept	0.13 (0.32)	0.40	.692	
	Algorithm disclosure	-0.86 (0.38)	-2.28	.023*	0.42 [0.20; 0.88]
	Anthropomorphic design	0.73 (0.38)	1.92	.055.	2.07 [0.99; 4.40]
Model 3	Intercept	1.10 (1.21)	0.91	.362	
	Algorithm disclosure	-3.71 (1.49)	-2.49	.013*	0.02 [0.00; 0.39]
	Anthropomorphic design	-1.85 (1.47)	-1.26	.207	0.16 [0.01; 2.62]
	Trust	0.22 (0.28)	-0.77	.439	0.81 [0.45; 1.38]
	Algorithm disclosure:Trust	0.66 (0.34)	-1.93	.053.	1.94 [1.01; 3.93]
	Anthropomorphic design:Trust	0.60 (0.34)	1.77	.077.	1.82 [0.95; 3.35]

Table 3. Results of the logistic regression models. *p<.05, **p<.01, ***p<.001

To facilitate the interpretation of the results, we calculated the probabilities predicted by the third model of the logistic regression for the different levels of the independent variables. The results are depicted in

Figure 3. We dichotomized trust by defining low trust as the mean of trust minus the standard deviation (4.06-1.33 = 2.73) and high trust as the mean of trust plus the standard deviation (4.06+1.33 = 5.39).

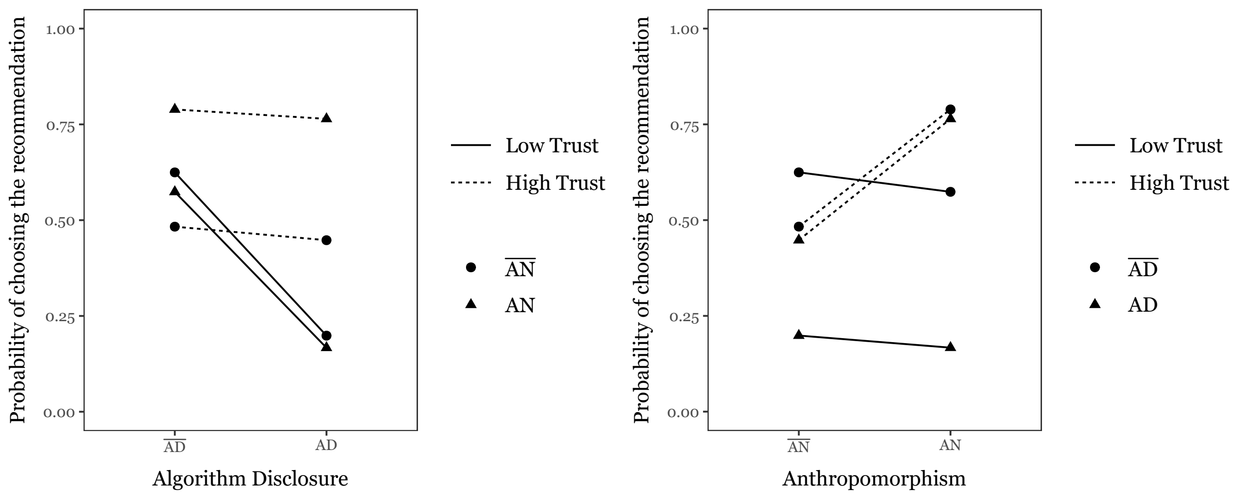

Figure 3. Results of the third logistic regression modeled as probabilities to choose the recommendation.

Discussion, limitations and future research

The adoption of AI-based recommendations in the job-seeking domain is a crucial determinant concerning the future success of artificial intelligence in HR. To shed light on the multi-faceted, complex issue of recommendation acceptance in the job-seeking domain, we conducted an empirical scenario-based vignette study using an online survey. We investigated the impact of algorithm aversion and the effects of anthropomorphic design features on the acceptance of AI-based job recommendations from a user's perspective. Our results highlight that recommendation acceptance in the job-seeking domain cannot be reduced to a small set of determinants, but is influenced by multiple factors. Our results suggest that algorithm aversion, triggered by algorithmic disclosure, is an influential factor on recommendation acceptance. In line with our first hypothesis, we found evidence that disclosing additional information on the algorithmic origin of the AI-based job recommendation in the treatment groups (AN_AD, $\overline{\text{AN}}$_AD) led to a significant decrease in the acceptance of the recommendation compared to the groups to which no additional information was disclosed (AN_$\overline{\text{AD}}$, $\overline{\text{AN}}$_$\overline{\text{AD}}$). With regard to our second hypothesis, we found a marginally significant effect of anthropomorphic design features on the acceptance of the job recommendation. This effect diminished, however, once we added the general trust in artificial intelligence in the job domain as a moderator to the model.

Our final model indicates that algorithmic disclosure has a significant, negative effect on the acceptance of AI-based recommendations in the job-seeking context, indicating that algorithm aversion is a highly influential and important factor in this domain. A more detailed analysis, including the marginally significant moderation effect of general trust in artificial intelligence in the job domain, showed that this effect loses strength with increasing general trust. Among low-trust individuals in the high algorithm disclosure groups (AN_AD, $\overline{\text{AN}}$_AD), the predicted probability of choosing the recommendation was approximately 20%, compared to roughly 60% among the low-trust individuals in the low algorithm disclosure groups ((AN_$\overline{\text{AD}}$, $\overline{\text{AN}}$_$\overline{\text{AD}}$). By contrast, high-trust individuals in the corresponding groups accepted the recommendation with a probability of roughly 50%, up to over 75%, with only minimal differences between the different algorithmic disclosure groups.

These results suggest that algorithm aversion can indeed hinder the acceptance of AI-based recommendations in high-stake contexts. Thus, our findings corroborate the findings of prior research on algorithm aversion (Burton et al. 2020; Castelo et al. 2019; Dietvorst et al. 2015) and strengthens the need for additional research in this field. Our findings regarding the moderating effect of general trust in the algorithm-using technology contribute to prior research by unveiling an influential factor that might partially drive algorithm aversion in high-stake decision contexts. Individuals with low trust seem to be more sensitive to algorithm disclosure and more prone to algorithm aversion. By contrast, disclosing information on individuals with high trust does not seem to induce algorithm aversion. To the best of our knowledge, this is the first study that reports such effects of general trust on algorithm aversion and, therefore, extends the research on trust in adoption behavior (Komiak and Benbasat 2008, 2006; Moussawi et al. 2020; Qiu and Benbasat 2009; Wang et al. 2016), and algorithm aversion. One should keep in mind, however, that the level of general trust in artificial intelligence could be related to the familiarity with such technology (Komiak and Benbasat 2008). Therefore, it might be the case that disclosing information on the algorithmic origin of the AI-based job recommendation did not lead to algorithm aversion in the high trust group, as they were already aware of this relation. Future research should thus investigate this more deeply, and take familiarity and the novelty of the disclosed information for the users into account.

With regard to our second hypothesis, our final regression did not show a significant main effect of anthropomorphic design features on the acceptance of AI-based job recommendations. A more detailed analysis, including the marginally significant moderation effect of general trust in artificial intelligence, showed that the effect of anthropomorphic design features depends on the level of general trust in artificial intelligence of the user. While for high-trust individuals in the anthropomorphism groups ($\mathrm{AN\_AD}$, $\mathrm{AN\_\overline{AD}}$) the predicted probability of choosing the recommendation increases to over 75% compared to around 45% in the no anthropomorphism groups (($\mathrm{\overline{AN}\_AD}$, $\mathrm{\overline{AN}\_\overline{AD}}$), for low trust individuals, the predicted probability to accept the recommendation in both anthropomorphism groups is roughly 2% - 5% below the probability in the no anthropomorphism groups.

These results suggest that anthropomorphic design features do not necessarily increase the acceptance of AI-based recommendations in the job-seeking context and are not in line with the majority of findings by prior research (Moussawi et al. 2020; Pak et al. 2012; Qiu and Benbasat 2009). We conjecture that in high-stake decision contexts, anthropomorphic design features might not be an effective measure to increase the acceptance of an AI-based recommendation agent's suggestions. If individuals are generally suspicious regarding the applicability of artificial intelligence in the job context and do not trust it, anthropomorphic design features do not contribute to an increased acceptance rate. It it conceivable that these individuals generally do not react to any persuasive approaches in these contexts, as they feel less ambivalence with regard to their rejective stance (Jonas et al. 2000; Zemborain and Johar 2007) and tend to ignore information that is not in line with their attitude (Rothman et al. 2017). Individuals with high trust and thus less suspicion in the anthropomorphism groups, on the other hand, might actively look for information confirming their prior attitude towards artificial intelligence in job-seeking (Rothman et al. 2017), such as positively perceived anthropomorphic design features (Epley et al. 2007; Qiu and Benbasat 2009; Wang et al. 2016). Further research is needed to investigate the underlying mechanisms driving the moderating effects of trust in this context.

Our results have multiple implications for academic research and contribute to the ongoing discussion regarding possible interventions to increase the acceptance of AI-based job recommendations.

First, prior studies in the research stream of recommender systems generated insights by outlining the impacts of different cognitive (Moussawi et al. 2020) and affective factors (Komiak and Benbasat 2008), thereby focusing on the consumer in commercial contexts. With the present study, we extend this research to the job-seeking context that is characterized by higher stakes involved in the decision. Our model shows that algorithm aversion and anthropomorphism affect the acceptance of AI-based systems. This emphasizes that technology acceptance in a high-stake context can be influenced by various factors that have not been considered in prior research so far.

Second, prior research on recommender systems has primarily investigated factors explaining the acceptance of recommendations by systems that rely on conventional information technology (Lu et al. 2015). With the increasing demand for and prevalence of artificial intelligence (Castelvecchi 2016), it is crucial to discuss if the factors influencing adoption behavior differ between recommendations based on conventional technology and AI-based recommendations. Our study is a first step in this direction, as we show that disclosing information on the algorithmic origin of a recommendation can lead to adverse effects on its acceptance in a context that is characterized by high personal stakes, thereby addressing affective factors. However, the moderating effects of trust in our study emphasize the need for an integrated view of cognitive and affective determinants of technology acceptance.

Third, algorithm aversion receives increasing attention in the light of the development in AI-based systems. It refers to a general tendency of individuals to prefer human forecasts and recommendations over algorithm-based ones. Our study is among the first to investigate algorithm aversion in the high-stakes context of job-seeking. By contrast, prior research mainly focused on numeric estimation tasks (Dietvorst et al. 2018; Logg et al. 2019) or contexts where the consequences of a wrong decision are negligible for the individual (Yeomans et al. 2019). Our findings suggest that algorithm aversion, elicited by disclosing information on the algorithmic origin of a recommendation, can be a critical factor inhibiting the acceptance of AI-based job recommendations in such contexts.

In addition, to the best of our knowledge, our study is the first to investigate the effect of anthropomorphic design features on the acceptance of AI-based recommendations in the job-seeking context. As these decisions are characterized by much higher stakes than typical decisions in the B2C context in which the use of anthropomorphic design features has been investigated before (Pak et al. 2012; Qiu and Benbasat 2009), our study contributes to prior research by testing whether anthropomorphic design features are also an effective measure in the high-stake context. This contribution is especially important in light of the fact that high-stake decisions are not yet discussed in the literature on recommender systems. It is reasonable, however, to assume that due to the advances in artificial intelligence and machine learning, the accuracy and performance of recommender systems will further increase (Logg et al. 2019). Therefore, such systems will be increasingly implemented in high-stake contexts, such as medical decision-making or financial decision-making (e.g., high-volume exchange traded funds). Therefore, it is important to assess measures that could be used in these contexts to increase the acceptance of recommendation systems in these domains, as they have the potential to support individuals in achieving better decision outcomes.

For practitioners, our research has multiple valuable implications regarding the introduction of an AI-based job recommender system. First, and on a general level, disclosing information on the algorithmic origin of the recommendation could lead to adverse reactions regarding the acceptance of the recommendation. Therefore, it might be beneficial not to disclose such information. If the information needs to be disclosed for transparency reasons or due to privacy guidelines, additional measures to increase general trust should be taken. One promising approach in this regard could be the inclusion of assurance seals on the recommendation system's website (Odom et al. 2002; Özpolat et al. 2013). Second, anthropomorphic design features can have a positive effect on the acceptance of AI-based job recommendations. As they are very easy to implement and do not involve high costs, it is recommended that practitioners include anthropomorphic design features to potentially increase the rate of acceptance of AI-based recommendations. Other positive effects of such anthropomorphizations could be increased loyalty and emotional connection of customers with the company (Araujo 2018; Guido and Peluso 2015).

Although our research provides valuable results for practice and academia, it comes with some limitations that future research should try to address. The first limitation is that our sample consisted of 121 indivuals from the USA and the UK. A culturally more diverse sample would increase the external validity of the study. Hence, in future work, we plan to further test our model in order to evaluate cultural generalization. Second, our findings are solely based on a survey with a hypothetical scenario, thus the participants' decision whether or not to follow the recommendation of the AI-based system did not have actual consequences in their real lives. In future projects, our aim is to conduct a field study where we plan to implement an AI-based recommender system in a company and to evaluate user acceptance in real-world decision contexts. Lastly, to further expand our research, we call on fellow researchers to contribute from the IS domain or related domains, such as Human-Computer-Interaction.

Conclusion

The findings of this study contribute to both academia and practice. Regarding the research question, we show that disclosing detailed information on the algorithmic origin of an AI-based recommendation can lead to algorithm aversion in a high-stake context like job search. As a result, individuals are more likely to reject the recommendation. At the same time, the results of our study indicate that the use of anthropomorphic design features to communicate an AI-based job recommendation can increase user acceptance.

Acknowledgements

This project is funded by the Adecco Stiftung "New Ways for Work and Social Life" and the Bavarian State Ministry of Science and the Arts, coordinated by the Bavarian Research Institute for Digital Transformation (bidt).

References

Adomavicius, G., Bockstedt, J. C., Curley, S. P., and Zhang, J. 2019. "Reducing Recommender System Biases: An Investigation of Rating Display Designs," *MIS Quarterly: Management Information Systems* (43:4), pp. 1321–1341.

Aggarwal, P., and McGill, A. L. 2007. "Is That Car Smiling at Me? Schema Congruity as a Basis for Evaluating Anthropomorphized Products," *Journal of Consumer Research* (34:4), pp. 468–479.

Araujo, T. 2018. "Living up to the Chatbot Hype: The Influence of Anthropomorphic Design Cues and Communicative Agency Framing on Conversational Agent and Company Perceptions," *Computers in Human Behavior* (85), pp. 183–189.

Bartneck, C., Kulić, D., Croft, E., and Zoghbi, S. 2009. "Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots," *International Journal of Social Robotics* (1:1), pp. 71–81.

Benjamini, Y., and Hochberg, Y. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society: Series B (Methodological)* (57:1), pp. 289–300.

Benlian, A., Klumpe, J., and Hinz, O. 2019. "Mitigating the Intrusive Effects of Smart Home Assistants by Using Anthropomorphic Design Features: A Multimethod Investigation," *Information Systems Journal* pp. 1–33.

Berg, J. M., Grant, A. M., and Johnson, V. 2010. "When Callings are Calling: Crafting Work and Leisure in Pursuit of Unanswered Occupational Callings," *Organization Science* (21:5), pp. 973–994.

Bodapati, A. V. 2008. "Recommendation Systems with Purchase Data," *Journal of Marketing Research* (45:1), pp. 77–93.

Bowling, N. A., Eschleman, K. J., and Wang, Q. 2010. "A Meta-analytic Examination of the Relationship Between Job Satisfaction and Subjective Well-Being," *Journal of Occupational and Organizational Psychology* (83:4), pp. 915–934.

Burton, J. W., Stein, M. K., and Jensen, T. B. 2020. "A Systematic Review of Algorithm Aversion in Augmented Decision Making," *Journal of Behavioral Decision Making* (33:2), pp. 220–239.

Castelo, N., Bos, M. W., and Lehmann, D. R. 2019. "Task-Dependent Algorithm Aversion," *Journal of Marketing Research* (56:5), pp. 809–825.

Castelvecchi, D. 2016. "Can We Open the Black Box of AI?" *Nature* (538:7623), pp. 20–23.

Cyr, D., Hassanein, K., Head, M., and Ivanov, A. 2007. "The Role of Social Presence in Establishing Loyalty in E-Service Environments," *Interacting with Computers* (19:1), pp. 43–56.

Cyr, D., Head, M. M., Larios, H., and Pan, B. 2009. "Exploring Human Images in Website Design: A Multi-Method Approach," *MIS Quarterly* (33:3), pp. 539–566.

Dana, J., and Thomas, R. 2006. "In Defense of Clinical Judgment … and Mechanical Prediction," *Journal of Behavioral Decision Making* (19:5), pp. 413–428.

Dawes, R. M. 1979. "The Robust Beauty of Improper Linear Models in Decision Making," *American Psychologist* (34:7), pp. 571–582.

Dawes, R. M., Faust, D., and Meehl, P. E. 1989. "Clinical Versus Actuarial Judgment," *Science* (243:4899), pp. 1668–1674.

Diab, D. L., Pui, S. Y., Yankelevich, M., and Highhouse, S. 2011. "Lay Perceptions of Selection Decision Aids in US and Non-US Samples," *International Journal of Selection and Assessment* (19:2), pp. 209–216.

Dietvorst, B. J., Simmons, J. P., and Massey, C. 2015. "Algorithm Aversion: People Erroneously Avoid Algo-

rithms After Seeing Them Err." *Journal of Experimental Psychology: General* (144:1), pp. 114–126.
Dietvorst, B. J., Simmons, J. P., and Massey, C. 2018. "Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them," *Management Science* (64:3), pp. 1155–1170.
Dijksterhuis, A., Bos, M. W., Nordgren, L. F., and van Baaren, R. B. 2006. "On Making the Right Choice: The Deliberation-Without-Attention Effect," *Science* (311:5763), pp. 1005–1007.
Duan, Y., Edwards, J. S., and Dwivedi, Y. K. 2019. "Artificial Intelligence for Decision Making in the Era of Big Data – Evolution, Challenges and Research Agenda," *International Journal of Information Management* (48), pp. 63–71.
Eastwood, J., Snook, B., and Luther, K. 2012. "What People Want From Their Professionals: Attitudes Toward Decision-making Strategies," *Journal of Behavioral Decision Making* (25:5), pp. 458–468.
Einhorn, H. J. 1986. "Accepting Error to Make Less Error," *Journal of Personality Assessment* (50:3), pp. 387–395.
Epley, N., Waytz, A., and Cacioppo, J. T. 2007. "On Seeing Human: A Three-Factor Theory of Anthropomorphism." *Psychological Review* (114:4), pp. 864–886.
Fehr-Duda, H., Bruhin, A., Epper, T., and Schubert, R. 2010. "Rationality on the Rise: Why Relative Risk Aversion Increases with Stake Size," *Journal of Risk and Uncertainty* (40:2), pp. 147–180.
Finch, J. 1987. "The Vignette Technique in Survey Research," *Sociology* (21:1), pp. 105–114.
Gomez-Uribe, C. A., and Hunt, N. 2015. "The Netflix Recommender System: Algorithms, Business Value, and Innovation," *ACM Transactions on Management Information Systems* (6:4).
Gong, L. 2008. "How Social is Social Responses to Computers? The Function of the Degree of Anthropomorphism in Computer Representations," *Computers in Human Behavior* (24:4), pp. 1494–1509.
Gruber, D., Aune, A., and Koutstaal, W. 2018. "Can Semi-Anthropomorphism Influence Trust and Compliance?" in *Proceedings of the Technology, Mind, and Society (TechMindSociety '18)*, , New York, New York, USA: ACM Press.
Gruber, D. S. 2018. *The Effects of Mid-range Visual Anthropomorphism on Human Trust and Performance Using a Navigation-based Automated Decision Aid*, Ph.D. thesis, University of Minnesota.
Guido, G., and Peluso, A. M. 2015. "Brand Anthropomorphism: Conceptualization, Measurement, and Impact on Brand Personality and Loyalty," *Journal of Brand Management* (22:1), pp. 1–19.
Hosmer, D. W., Lemeshow, S., and Sturdivant, R. X. 2013. *Applied Logistic Regression*, New Jersey, USA: John Wiley & Sons, 3rd ed.
Hu, R., and Pu, P. 2009. "Acceptance issues of personality-based recommender systems," *RecSys'09 - Proceedings of the 3rd ACM Conference on Recommender Systems* pp. 221–224.
Jonas, K., Broemer, P., and Diehl, M. 2000. "Attitudinal Ambivalence," *European Review of Social Psychology* (11:1), pp. 35–74.
Khasawneh, O. Y. 2018. "Technophobia without Boarders: The Influence of Technophobia and Emotional Intelligence on Technology Acceptance and the Moderating Influence of Organizational Climate," *Computers in Human Behavior* (88), pp. 210–218.
Komiak, S., and Benbasat, I. 2008. "A Two-Process View of Trust and Distrust Building in Recommendation Agents: A Process-Tracing Study," *Journal of the Association for Information Systems* (9:12), pp. 727–747.
Komiak, S. Y. X., and Benbasat, I. 2006. "The Effects of Personalization and Familiarity on Trust and Adoption of Recommendation Agents," *MIS Quarterly* (30:4), pp. 941–960.
Laumer, S., Gubler, F., Maier, C., and Weitzel, T. 2018. "Job Seekers' Acceptance of Job Recommender Systems: Results of an Empirical Study," *Proceedings of the 51st Hawaii International Conference on System Sciences* pp. 3914–3923.
LeCun, Y., Bengio, Y., and Hinton, G. 2015. "Deep learning," *Nature* (521:7553), pp. 436–444.
Logg, J. M., Minson, J. A., and Moore, D. A. 2019. "Algorithm Appreciation: People Prefer Algorithmic to Human Judgment," *Organizational Behavior and Human Decision Processes* (151), pp. 90–103.
Longoni, C., Bonezzi, A., and Morewedge, C. K. 2019. "Resistance to Medical Artificial Intelligence," *Journal of Consumer Research* (46:4), pp. 629–650.
Lu, J., Wu, D., Mao, M., Wang, W., and Zhang, G. 2015. "Recommender System Application Developments: A Survey," *Decision Support Systems* (74), pp. 12–32.
Malinowski, J., Wendt, O., Keim, T., and Weitzel, T. 2006. "Matching People and Jobs: A Bilateral Rec-

ommendation Approach," *Proceedings of the 39th Hawaii International Conference on System Sciences (HICSS)* pp. 1–9.
Meehl, P. E. 1954. *Clinical versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence.*, Minneapolis: University of Minnesota Press.
Mori, M., MacDorman, K., and Kageki, N. 2012. "The Uncanny Valley [From the Field]," *IEEE Robotics & Automation Magazine* (19:2), pp. 98–100.
Mourey, J. A., Olson, J. G., and Yoon, C. 2017. "Products as Pals: Engaging with Anthropomorphic Products Mitigates the Effects of Social Exclusion," *Journal of Consumer Research* (44:2), pp. 414–431.
Moussawi, S., Koufaris, M., and Benbunan-Fich, R. 2020. "How Perceptions of Intelligence and Anthropomorphism Affect Adoption of Personal Intelligent Agents," *Electronic Markets* .
Odom, M. D., Kumar, A., and Saunders, L. 2002. "Web Assurance Seals: How and Why They Influence Consumers' Decisions," *Journal of Information Systems* (16:2), pp. 231–250.
Özpolat, K., Gao, G. G., Jank, W., and Viswanathan, S. 2013. "The Value of Third-Party Assurance Seals in Online Retailing: An Empirical Investigation," *Information Systems Research* (24:4), pp. 1100–1111.
Pak, R., Fink, N., Price, M., Bass, B., and Sturre, L. 2012. "Decision Support Aids with Anthropomorphic Characteristics Influence Trust and Performance in Younger and Older Adults," *Ergonomics* (55:9), pp. 1059–1072.
Palan, S., and Schitter, C. 2018. "Prolific.ac—A Subject Pool for Online Experiments," *Journal of Behavioral and Experimental Finance* (17), pp. 22–27.
Park, D. H., Kim, H. K., Choi, I. Y., and Kim, J. K. 2012. "A Literature Review and Classification of Recommender Systems Research," *Expert Systems with Applications* (39:11), pp. 10,059–10,072.
Peer, E., Brandimarte, L., Samat, S., and Acquisti, A. 2017. "Beyond the Turk: Alternative Platforms for Crowdsourcing Behavioral Research," *Journal of Experimental Social Psychology* (70), pp. 153–163.
Pfeuffer, N., Benlian, A., Gimpel, H., and Hinz, O. 2019. "Anthropomorphic Information Systems," *Business & Information Systems Engineering* (61:4), pp. 523–533.
Promberger, M., and Baron, J. 2006. "Do Patients Trust Computers?" *Journal of Behavioral Decision Making* (19:5), pp. 455–468.
Pryce-Jones, J. 2010. *Happiness at Work: Maximizing Your Psychological Capital for Success.*, Oxford, UK: Wiley-Blackwell.
Qiu, L., and Benbasat, I. 2009. "Evaluating Anthropomorphic Product Recommendation Agents: A Social Relationship Perspective to Designing Information Systems," *Journal of Management Information Systems* (25:4), pp. 145–182.
Resnick, P., Varian, H. R., and Editors, G. 1997. "Recommender Systems," *Communications of the ACM* (40:3), pp. 56–58.
Riegelsberger, J., Sasse, M. A., and McCarthy, J. D. 2003. "Shiny Happy People Building Trust?" in *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2003)*, , New York, New York, USA: ACM Press.
Rothman, N. B., Pratt, M. G., Rees, L., and Vogus, T. J. 2017. "Understanding the Dual Nature of Ambivalence: Why and When Ambivalence Leads to Good and Bad Outcomes," *Academy of Management Annals* (11:1), pp. 33–72.
Schultze, U., and Brooks, J. A. M. 2019. "An Interactional View of Social Presence: Making the Virtual Other "Real"," *Information Systems Journal* (29:3), pp. 707–737.
Sharma, A., Hofman, J. M., and Watts, D. J. 2015. "Estimating the Causal Impact of Recommendation Systems from Observational Data," *Proceedings of the 16th ACM Conference on Economics and Computation (EC'15)* pp. 453–470.
Short, J., Williams, E., and Christie, B. 1976. *The Social Psychology of Telecommunications*, London: Wiley.
Sinkovics, R. R., Stöttinger, B., Schlegelmilch, B. B., and Ram, S. 2002. "Reluctance to Use Technology-related Products: Development of a Technophobia Scale," *Thunderbird International Business Review* (44:4), pp. 477–494.
Steinbrück, U., Schaumburg, H., Duda, S., and Krüger, T. 2002. "A Picture Says More Than a Thousand Words," in *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2002)*, , New York, New York, USA: ACM Press.
van Esch, P., Black, J. S., and Ferolie, J. 2019. "Marketing AI Recruitment: The Next Phase in Job Application and Selection," *Computers in Human Behavior* (90), pp. 215–222.

Wanberg, C., Zhu, J., and Van Hooft, E. 2010. "The Job Search Grind: Perceived Progress, Self-Reactions, and Self-Regulation of Search Effort," *Academy of Management Journal* (53:4), pp. 788–807.
Wang, W., Qiu, L., Kim, D., and Benbasat, I. 2016. "Effects of Rational and Social Appeals of Online Recommendation Agents on Cognition- and Affect-basedTtrust," *Decision Support Systems* (86), pp. 48–60.
Yeomans, M., Shah, A., Mullainathan, S., and Kleinberg, J. 2019. "Making Sense of Recommendations," *Journal of Behavioral Decision Making* (32:4), pp. 403–414.
Zemborain, M., and Johar, G. 2007. "Attitudinal Ambivalence and Openness to Persuasion: A Framework for Interpersonal Influence," *Journal of Consumer Research* (33:4), pp. 506–514.