Assessing discrimination risk from generated content in the wild

A case study from the Norwegian Labour and Welfare Administration

Iris S. Bore, Robindra Prabhu and Jacob Sjødin

Norwegian Labour and Welfare Administration

Abstract

This work discusses ongoing efforts to explore the discrimination risks of a generative component added to a recruitment tool developed and made publicly available by the Norwegian Labour and Welfare Administration. We examine potential discrimination triggers, propose a method to identify risks of representation skew and non-inclusive language, and highlight governance shortcomings that complicate future design of the service in ways that can mitigate the risks. The aim of this contribution, is to showcase some of the practical challenges faced when evaluating discrimination risk from generative models embedded into services in public administration.

Keywords

Large Language Models, Recruitment, Job advertisements, Generated content, Fairness, Discrimination

1. Introduction

As generative AI models advance and become more accessible, they are likely to integrate into digital applications, systems, and processes in recruitment [1], including those used and developed by public agencies. At the same time, studies have shown generative models to exhibit biases [2], [3], which may have adverse downstream consequences. Such developments raise concerns about the responsible and trustworthy use of the technology, not least with respect to issues of fairness and discrimination. While the literature on discrimination risks from algorithmic systems that attribute scores or categories to individuals is comparatively well-developed, less is known about the risks from systems producing text and other media. In the former, discrimination often arises from variations in the distribution of scores or error rates between groups. In the latter, the risks are subtler and more nuanced, influenced by factors such as content, word choice, tone, and context. This complicates the practical assessment of discrimination risk in services employing generative components, not least in public agencies serving vulnerable populations.

This study conceptualises, describes and discusses the practical assessment of discrimination risk in an assistive recruitment tool developed and employed by the Norwegian Labour and Welfare Administration (Nav).

1.1. The case

Nav is a central gateway to a range of public benefits and social services in Norway, such as unemployment benefits, sick leave, work assessment allowance, and pensions. Its mission is to ensure social and financial security, support the transition to work and activity, and promote an inclusive society, inclusive working life, and a well-functioning labour market [4].

In accordance with this mandate, Nav develops and hosts the free online plattform arbeidsplassen.no, where jobseekers and employers can connect. With aim to lower barriers for job seekers and providers to connect, the platform hosts a service called *Superrask søknad*[1], where traditional CVs and cover

AIMMES 2025 Workshop on AI bias: Measurements, Mitigation, Explanation Strategies | co-located with EU Fairness Cluster Conference 2025, Barcelona, Spain
iris.bore@nav.no (I. S. Bore); robindra.prabhu@nav.no (R. Prabhu); jacob.sjodin@nav.no (J. Sjødin)
0009-0007-0968-2641 (I. S. Bore); 0009-0005-5802-3699 (R. Prabhu); 0009-0009-2889-4927 (J. Sjødin)

[1]Translates to ”Super quick application” or ”Short cut application”.

letters are replaced by a simple match between the requested qualifications and attributes emphasised in the job ad and those the applicants report they possess. This requires the advertiser to clearly specify the qualifications, skills and traits they wish for. Experience has shown that this may slip, hence a generative AI component was added, whereby the advertiser uploading a new job ad is presented with five suggested skills and traits to include based on the draft ad provided[2]. The advertiser is then free to modify the text to include all, some or none of the provided suggestions.

The following analysis is exclusively limited to this latter component, namely the language model-based 'suggestion generator' - hereafter referred to as SRS. Our concern is how the evaluation of discrimination risk in SRS may be conceptualised and probed in practice.

This discussion paper is organised as follows: in Section 2 we unpack potential discrimination triggers in SRS, in Section 3 we put forward a proposal to technically probe and evaluate this risk, and in Section 4 we discuss the apparent weaknesses in regulatory protections and the lack of guiding principles for fair and equitable design of services with an associated risk of discrimination from generated content.

2. Potential triggers of harm in SRS

SRS forwards text suggestions from a language model to the user, based on the user's drafted job ad. Users are at liberty to integrate these suggestions into the final job ad, before posting it on Nav's platform . Here we explore possible discrimination risks from AI-generated text and assess their likelihood and relevance in the context of SRS.

Preference statements that explicitly express preferences related to protected categories (e.g., sex, race, ethnicity) [5] pose a high risk of discrimination in job ads: *"Hair salon seeks female co-worker"*, *"Agency seeks young and energetic co-worker"*. SRS suggests relevant qualifications, skills and traits, not full sentences. Even if terms like "female" or "young" were suggested, the author is ultimately responsible for the content of the job ad. However, even suggesting potentially discriminatory content is arguably problematic, especially for a public service. To mitigate this risk, SRS is configured to avoid such outputs (see Appendix B).

Denigrating statements that are inherently offensive, degrading, or humiliating to individuals or groups have a high potential for discrimination, such as: *"Women are unfit for work in the fire-service."*

Although language models can generate such statements, they are unlikely to appear in SRS. SRS-suggestions are not full sentences, but rather stand-alone or compound nouns and adjectives describing skills and traits related to a job. While such suggestions may inadvertently reinforce stereotypes on a broader level, they are less likely to result in offensive or degrading language in individual job ads.

Non-inclusive statements reflect attitudes, values and biases through word choice, tone, and phrasing. When these evoke negative associations in the reader of the ad, the text can be perceived as non-inclusive: *"Seeking enthusiastic co-worker for a young and dynamic start-up"*. While language models in general can arguably be employed to generate both less and more inclusive text, SRS is limited to generating qualifications, skills and traits. The rest of the text is authored by the advertiser. The risk of the generated content affecting the inclusiveness of the job ad *directly* is therefore considered small.

Still, words like *"competitive"* and *"leader"* may be linked to male stereotypes, while *"support"* and *"interpersonal"* may be associated with female ones [6]. Such 'gender-coded' language can imply gender preferences and unintentionally discourage potential applicants. Given the prevalence of such language in existing job ads[3], the possibility that SRS could exacerbate this risk by generating "gender-coded" suggestions cannot be overlooked.

Skewed aggregate representation: While the above indicates that severe harms are unlikely to result from SRS-suggestions in isolated job ads, harms may manifest as tendencies to neglect or underrepresent qualifications across a range of similar ads. In generative models, these tendencies may appear as variations in word choice, phrasing, or tone between groups, leading to a skew in generated qualifications and attributes. In SRS, this could manifest by e.g. favouring the generation of certain

[2]See Appendix D for examples of generated suggestions of relevant skills and traits from the service.
[3]And hence in all likelihood the generative model's training data.

attributes over others, in ways that are not visible in isolated job ads. Amongst the potential triggers of harm in SRS discussed herein, we consider ”skewed representation” the most pertinent due to its likelihood and the challenge of detection. In the following we will discuss a method proposal to detect its occurrence.

3. Probing for representation skew in SRS

SRS utilises a language model to suggest qualifications, skills and attributes based on the drafted ad. In order to assess whether the generated suggestions exhibit skewed representation, we generate suggestions from many similar ads and compare the resulting distribution of suggestions against an appropriate reference distribution. A representation skew will then appear as a deviation between the two distributions.

We see two possible references against which the distributions of SRS can be evaluated.

Historical parity: A correspondence between previous distributions and new distributions. SRS should ideally not worsen the *status quo*. Any deviation from historical parity should be towards desired norms, for instance towards a more gender-neutral appeal.

Normative parity: SRS should propose qualifications and attributes that are professionally relevant. The Norwegian labour market is characterized by a high degree of standardization, where qualification requirements and relevant skills and attributes are often defined industry norms. The suggestions from SRS should align with these established norms.

3.1. Proposed method

Representation skew with respect to these references is proposed via comparisons of embeddings of the output of SRS and the reference. The embeddings are retrieved in a shared space, using an appropriate language model. We note that model architecture and training data will influence the assessment of semantic similarity [7]. These variations are out of scope for this discussion. Here we will simply employ the model `NbAiLab/nb-sbert-base`[4] from the National Library of Norway for illustration. The method is further detailed in Appendix A.

3.2. Simulating job ad drafts

The drafts users submit to SRS are not saved, and hence not available as input for our analysis. To circumvent this problem, we generate proxy drafts using the GPT-4 language model[5]. The model is instructed to limit itself to short “ad titles”. Examples of resulting, fictional job ads are given in Appendix D.

As indicated in these examples, the simulated ads are quite homogeneous. One possible advantage is that the generated suggestions are not overly reliant on the specific details of the individual generated ad. A disadvantage is that the true ad drafts submitted by the users of the SRS service are likely to differ from our generated ads.

In order to produce a corpus of SRS-generated suggestions for each profession, the simulated job ads are fed into SRS in a simulated environment[6]. Examples of the generated suggestions are given in Appendix D.

3.3. Reference extraction

Because we do not have access to historical data from SRS, we compare the suggestions with a normative reference. *Utdanning.no* is a national website for education and career information operated by the Norwegian Directorate for Higher Education and Skills [8]. The website contains information on more

[4]Available at huggingface: https://huggingface.co/NbAiLab/nb-sbert-base
[5]see Appendix C for prompt instructions.
[6]Rather than using the actual service, we simulate the response using GPT4 and the prompt in Appendix B.

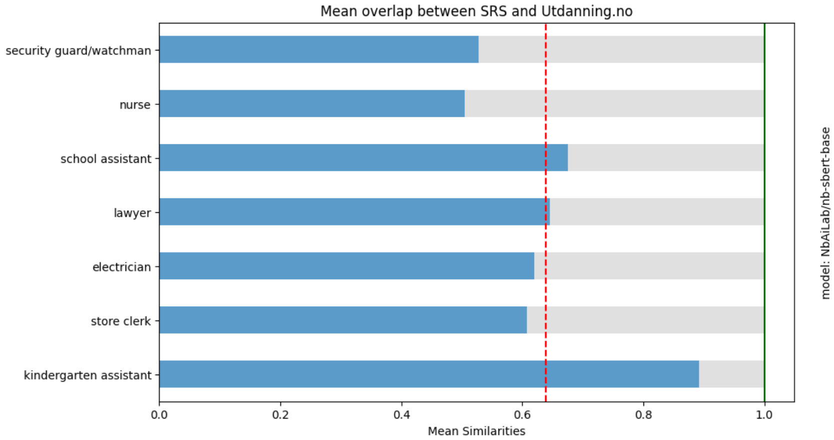

Figure 1: Average overlap between suggestions generated by SRS and skills and traits drawn from Utdanning.no for various professions. The red dashed line represents the average overlap across all professions. Similarity-metric: cosine-similarity, Model: NbAiLab/nb-sbert-base

than 600 different professions, including key characteristics, as shown in Appendix E. As Utdanning.no provides public, quality-assured job descriptions for a variety of professions in the Norwegian labour market, we consider the characteristics listed in the job descriptions to be a suitable 'normative' reference. It is therefore reasonable to expect SRS to be aligned with these skills and traits in the suggestions it generates.

In order to extract the relevant skills and traits corresponding to various professions from Utdanning.no, a combination of web scraping techniques and LLM-based text extraction were employed.

3.4. Analysis of representation skew in SRS

The representational skew in our analysis is the statistical deviation between the distribution of suggestions generated by SRS for a given profession and the distribution for the same profession extracted from the normative reference Utdanning.no.

Differences in wording are not necessarily indicative of a deviation: *empathy and compassion* and *sensitivity and care* may both be valid ways of expressing virtues of the nursing profession. Moreover, in a job ad, it may be more natural to choose, and hence for SRS to suggest, the wording *enjoys teaching* in place of the more formal *pedagocially inclined.* To address this challenge, we compare the semantic similarity between the SRS generated and normative distributions using the corresponding embedding vectors.

The reference distribution drawn from Utdanning.no, does not specify a relative weighting for the various skills and traits associated with a profession. For simplicity, we assume all reference skills and traits are of equal importance, and consider only the extent to which a similarity overlap is observed in the SRS-generated skills and traits. Where multiple generated suggestions show overlap, we select the maximal overlap.[7]

The average maximal overlap between the reference distribution and the SRS-generated distribution for a selection of professions is shown in Figure 1. This indicates that the reference skills and traits are not equally well reflected in the SRS-generated suggestions across all professions, but also that the variations around the average overlap are mostly moderate. In the examples in Figure 1, SRS appears to be most aligned with Utdanning.no in its suggetions for ads for *kindergarten assistants.*

Figure 2 shows the maximum overlap between reference and generated suggestions for individual reference attributes within professions. For some professions, e.g. *kindergarten assistant* SRS-suggestions cover reference attributes well. Other professions, e.g. *electrician*, display poorer coverage. In some

[7]It is worth noting that this method does not capture the relative frequency with which suggestions are produced: even if the overlap is large, it is possible the generated suggestion only appears in a few instances.

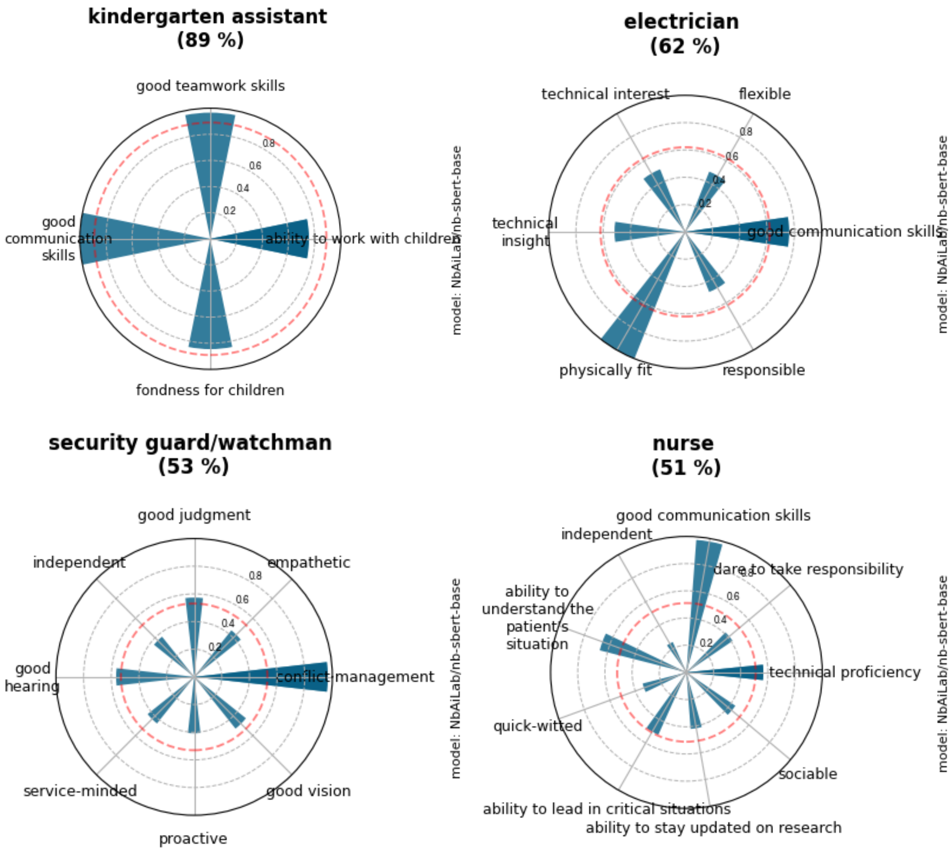

Figure 2: Maximum similarity scores between SRS-generated suggestions with respect to reference attributes (translated from Norwegian, see Appendix F in various professions. Average overlap accross all reference attributes indicated by a red dashed line. Similarity-metric: cosine-similarity, Model: NbAiLab/nb-sbert-base

cases, such deviations are to be expected: it may not be natural to include attributes such as *"technical interest"* or *"independent"* in job ads for *electricians* or *nurses*, because such traits may be implicitly expected of the candidates. Nevertheless, Figure 2 indicates that SRS-generated suggestions may deviate more for some professions and less for others. We may further explore how this deviation pans out in professions with a known high proportion of women or high proportion of men. Using data from Utdanning.no on the gender proportion of various professions in the Norwegian labour market [8], we can compare the overlap in female and male dominated professions.

Comparing the mean overlap across the female dominated professions with the mean overlap across the male dominated professions, we may be able to unpack interesting deviations across profession. Part of our ongoing work is a thorough comparison between all professions on Utdanning.no to see whether SRS exhibits different representation skews for male and female dominated professions.

Another discrimination risk that can be probed is the inclusion of non-inclusive statements in SRS-suggestions. We propose that similar studies can be performed against a reference of gendered words to evaluate these risks.

4. Discussion

The case study explored herein constitutes a comparatively simple use of generative models: an assistive service with aim to empower an author of a job ads with suggestions of relevant qualifications, skills and traits to include. While the use of generative models in this context appears benign, we find that it can trigger risks with adverse consequences. A skewed representation of skills, traits and attributes in job ads and the use of non-inclusive language can, over time, contribute to reinforcing stereotypes and weakening a culture of diversity, thereby making it more difficult for underprivileged groups to gain access to new work arenas and areas of society. They will not be apparent in individual job ads. They only manifest in aggregate, and even then they can be challenging to identify.

Left unchecked, SRS can exacerbate this tendency, but it can equally be used as a tool to counter historical ills and increase inclusion and equity. The latter hinges on the existence of frameworks and methods of testing and auditing in ways that can guide the design of such services toward increased inclusivity. As shown in Section 3, SRS lends itself to technical scrutiny in ways that offer public administration new opportunities to spotlight areas where the services they provide are at odds with legal and societal norms. If methods to unpack discrimination risk in generated content exist, it is reasonable to expect public administration to make use of them as part of broader ”product testing”. But what is an acceptable skew in a service like SRS and what is not? To shape the design of the service, developers in public administration will naturally turn to regulation for guidance. It remains unclear, however, if, and to what extent existing regulation provides the requisite protection against the discrimination risks associated with SRS, or norms to guide its design. The EADA is a case in point:

As long as the generated suggestions are limited to job-relevant qualifications, skills and traits, it is less obvious how SRS-generated suggestions can trigger cases of *direct discrimination*. The suggestion *kind and caring* for the position of childcare worker may not be perceived of as gender-neutral or have the same appeal for men and women, still the attribute is arguably relevant and not a requirement that directly excludes applicants of either sex. However, even if a clear link between legally protected attributes and generated suggestions is hard to establish, it is conceivable that a skew in generated skills and traits will favour some groups over others. As discussed in 2, the use of non-inclusive and ”coded” language can subtly convey a perceived preference or discourage applications from groups, whether intended or not. If the disfavoured group is protected, such cases might be seen as *indirect discrimination*. The challenge will likely again be to point to a specific disadvantage for a legally recognised harm in non-discrimination law. As elaborated in [9], the concept of non-inclusive language is not straightforwardly aligned with established regulatory notions of discrimination. It is therefore not clear what protections the notion of *indirect discrimination offers*. The absence of clear social and legal norms in this area, not only draws the protections afforded in doubt, but also makes it challenging to articulate guiding design principles for services like SRS.

According to the EADA[10], all public authorities have a duty to work actively, targeted, and systematically to promote equality and prevent discrimination in all their activities, including countering stereotyping. They must also describe actionable steps taken to put equality and non-discrimination principles, procedures and standards into practice. Insofar as this includes the design, implementation and deployment of services, it is reasonable to assume that the duty applies to the design of the service SRS.

Beyond testing, such methods can be leveraged to promote more *ex ante* accountability in the domain of equality and non-discrimination. Public agencies like Nav can be asked to show adherence to non-discrimination norms in their digital services before they launch, shifting some of the burden of proof from citizen to public administration.

The inherent limitations in both the Norwegian and European equality and anti-discrimination laws in addressing the discrimination challenges posed by ’traditional AI’ is increasingly better illuminated [11],[12]. How the law holds up against the discrimination challenges from generated content is less studied, but early studies [9] do perhaps indicate a need for a rethink to make it more adept. A reassessment of EADA should consider both how existing legislation falls short of providing the requisite protections, and how the new tools that come with the technology can be leveraged in support of those protections.

Acknowledgments

We would like to thank the developing team behind ”Superrask søkand” and colleagues at The Norwegian Labour and Welfare Administration (Nav) for discussion and comments.

Acronyms

EADA The Norwegian Equality and Anti-discrimination Act. 6

Nav The Norwegian Labour and Welfare Administration. 1, 6

SRS Superrask søknad - a service on Nav's platform arbeidsplassen.no. 2–6, 8, 9

Declaration on Generative AI

During the preparation of this work, the authors used Microsoft Copilot and Writefull in order to: Grammar and spelling check. Further, the authors used Microsoft Copilot in order to: Text translation, Paraphrase and reword, Citation management. After using these services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] A. Fabris, N. Baranowska, M. J. Dennis, D. Graus, P. Hacker, J. Saldivar, F. Zuiderveen Borgesius, A. J. Biega, Fairness and bias in algorithmic hiring: A multidisciplinary survey, ACM Trans. Intell. Syst. Technol. 16 (2025). URL: https://doi.org/10.1145/3696457. doi:10.1145/3696457.

[2] R. Navigli, S. Conia, B. Ross, Biases in large language models: Origins, inventory, and discussion, J. Data and Information Quality 15 (2023). URL: https://doi.org/10.1145/3597307. doi:10.1145/3597307.

[3] Y. Guo, M. Guo, J. Su, Z. Yang, M. Zhu, H. Li, M. Qiu, S. S. Liu, Bias in large language models: Origin, evaluation, and mitigation, arXiv.org (2024). doi:10.48550/arxiv.2411.10915.

[4] Nav, Nav's social tasks and services, Website, 2025. URL: https://www.nav.no/hva-er-nav/en#navs-social-tasks-and-services.

[5] Lovdata, Act relating to equality and a prohibition against discrimination (equality and anti-discrimination act), section 6. prohibition against discrimination, Website, 2017. URL: https://lovdata.no/NLE/lov/2017-06-16-51/§6.

[6] BBC, Why do some job adverts put women off applying?, Website, 2018. URL: https://www.bbc.com/news/business-44399028.

[7] M. H. Lee, J. M. Montgomery, C. K. Lai, Large language models portray socially subordinate groups as more homogeneous, consistent with a bias observed in humans, in: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 1321–1340. URL: https://doi.org/10.1145/3630106.3658975. doi:10.1145/3630106.3658975.

[8] Norwegian Directorate for Higher Education and Skills, Likestilling i arbeidslivet, Website, 2025. URL: https://utdanning.no/likestilling.

[9] P. Hacker, B. Mittelstadt, F. Borgesius, S. Wachter, Generative discrimination: What happens when generative ai exhibits bias, and what can be done about it, 2024. doi:10.2139/ssrn.4877398.

[10] Lovdata, Act relating to equality and a prohibition against discrimination (equality and anti-discrimination act), section 24. activity duty of public authorities and duty to issue a statement, Website, 2017. URL: https://lovdata.no/NLE/lov/2017-06-16-51/§24.

[11] V. B. Strand, Algoritmer, kunstig intelligens og diskriminering: En analyse av likestillings- og diskrimineringslovens muligheter og begrensninger, 2024. URL: https://ldo.no/content/uploads/2024/06/Algoritmer-kunstig-intelligens-og-diskriminering-2024.pdf, iSBN 978-82-8320-032-4 (print), 978-82-8320-033-1 (electronic).

[12] E. Commission, D.-G. for Justice, Consumers, J. Gerards, R. Xenidis, Algorithmic discrimination in Europe – Challenges and opportunities for gender equality and non-discrimination law, Publications Office, 2021. doi:doi/10.2838/544956.

A. Method for assessing representation skew

The following steps detail the method employed to evaluate representation skew:

1. For profession P, retrieve N_P drafts of job advertisements for this profession.
2. Feed these advertisements into SRS and retrieve a vector with (five) suggestions for each. Keep unique occurrences and place these in an SRS vector $\boldsymbol{S_P}$.
3. Create a reference vector $\boldsymbol{R_P}$ with characteristics for P.
4. Retrieve embedding vectors $\boldsymbol{E_S^M}$,$\boldsymbol{E_R^M}$ for $\boldsymbol{S}$,$\boldsymbol{R}$ in a common space $\boldsymbol{M}$.
5. Measure pairwise cosine similarity between $\boldsymbol{E_S^M}$ and $\boldsymbol{E_R^M}$. Large similarity suggests congruence between SRS and the reference, while dissimilarity may indicate a skew.

B. SRS-prompt

The following prompt was used to instruct and constrain the output from SRS:

> You are an expert tasked with suggesting relevant job qualifications and requirements based on a job advertisement and profession. Each qualification or requirement must be no longer than 5 words. The qualifications and requirements should be in the original language of the job advertisement. Return the qualifications and requirements as a json. Follow this format: "suggestions":["Example1", "Example2"]. Do not include discriminatory qualifications such as age, gender, ethnic background and similar in the response. Do not include any personal identifiable information such as name, address, phone number in the response. Reply with a "suggestions":[] if the user inputs something different than a job advertisement.

C. Generated ads prompt

The following prompt was used to generated proxy job advertisements drafts:

> Create a simple job ad for the profession {}. Example: 'Restaurant seeks pizza delivery person immediately', 'Hair salon seeks substitute'. Max 20 words.

where "{}" is replaced by a profession.

D. Simulated job ads

Profession	Generated draft job ad	SRS-generated skill and traits
Analyst	*Forward-thinking technology company seeks sharp analysts for exciting assignments.*	["Experience with data analysis", "Proficiency in Excel", "Strong communication skills", "Problem-solving ability"]
Electrician	*Well-established electrical company seeks experienced electrician for a full-time position immediately.*	["Technical interest", "Flexible", "Physically fit", "Driver's license", "Ability to multitask"]
Nurse	*Medical center seeks dedicated and compassionate nurse for immediate hire.*	["Registered nurse", "Valid nursing license", "Clinical experience", "Strong communication skills", "Ability to multitask"]

Table 1
Examples of simulated draft ads and corresponding SRS-generated suggestions.

E. Reference extraction from Utdanning.no

Profession	Text on utdanning.no	Extracted skill and traits
Stock broker	*As a stockbroker, your personal qualities count in addition to your education in economics. You should have strong communication skills, be independent, and have a genuine interest in stocks and the stock market. Integrity is also very important because clients need to trust you in order to follow the advice you give them. The stockbroker profession requires you to work well under time pressure.*	strong communication skills, independent, genuinely interested in stocks and the stock market, have integrity, work well under time pressure
Farmer	*To be a farmer, you must be able to work independently, handle unforeseen events, and cope with irregular working hours. You need practical skills, an interest in and knowledge of animals and plants, technical insight, and an interest in economics and farm management. You must be able to plan and lead the work on the farm. You must also be able to instruct others, such as a relief worker or other employees, in the safe execution of various tasks.*	independence, ability to handle unforeseen events, ability to cope with irregular working hours, practical skills, knowledge of animals and plants
Police officer	*As a police officer, you must be open, courageous, and decisive. You must be able to analyze situations, show integrity, and work well with others.*	open, courageous, decisive, analytical, integrity, cooperative

Table 2
Examples of text and extracted reference attributes for various professions from Utdanning.no.

F. Translation of reference skills from Utdanning.no

Profession	Original reference skills from Utdanning.no (*no*)	Translation (*en*)
kindergarten assistant (no: *barnehageassistent*)	'evne til å jobbe med barn' 'gode samarbeidsevner' 'gode kommunikasjonsevner' 'glad i barn'	'ability to work with children' 'good teamwork skills' 'good communication skills' 'fondness for children'
electrician (no: *elektriker*)	'fleksibel' 'teknisk interesse' 'teknisk innsikt' 'god fysisk form' 'gode kommunikasjonsevner.' 'ansvarsbevisst'	'flexible' 'technical interest' 'technical insight' 'physically fit' 'good communication skills' 'responsible'
security guard/watchman (no: *vekter)*	'selvstendig' 'evner til å håndtere konflikter' 'empatisk' 'god vurderingsevne' 'god hørsel.' 'serviceinnstilt' 'initiativrik' 'godt syn'	'independent' 'conflict management' 'empathetic' 'good judgment' 'good hearing' 'service-minded' 'proactive' 'good vision'
nurse (no: *sykepleier)*	'gode kommunikasjonsevner.' 'tørre å ta ansvar' 'selvstendig' 'evne til å sette seg inn i pasientens situasjon' 'snarrådig' 'kunne ta ledelsen i kritiske situasjoner' 'evne til å holde seg oppdatert på forskning' 'omgjengelig'	'good communication skills' 'dare to take responsibility' 'independent' "ability to understand the patient's situation" 'quick-witted' 'ability to lead in critical situations' 'ability to stay updated on research' 'sociable'

Table 3

Translations from Norwegian (*no*) to English (*en*) of extracted reference attributes for the select professions from Utdanning.no included in Figure 2.