Legal considerations when using AI in recruiting

impress

Contents

Disclaimer

The information provided on this whitepaper does not, and is not intended to, constitute legal advice and should not be relied on, nor treated as a substitute for specific advice relevant to particular circumstances. impress.ai is not a law firm, or otherwise engaged in the practice of law/ provision of legal advice. impress.ai does not accept any responsibility for any damages which may arise from reliance on information or materials published on this whitepaper. All information, content, and materials available in this paper are for general informational purposes only. Please direct all substantial legal queries to independent legal counsel.

Legal considerations when using AI in recruiting

Introductory remarks

The past decade has seen an exponential rise in automated decision-making software employed by recruiters at organizations around the world. A novel and mercurial legal consideration, Artificial Intelligence (AI) has little to no direct regulation globally, leaving the applicable laws disaggregated. Legal best practice standards are hence drawn from an adaptation of applicable privacy and data protection laws and employment law standards. This white paper is written for industry professionals looking for practical advice on legal considerations when looking to leverage AI recruitment software for their organizations.

AI can hold significant benefits for recruiters in comparison to the traditional process. With data-driven algorithms being able to collect, process and analyze candidate data with greater accuracy and efficiency, AI seemingly eliminated the bias-related issues that arise from human evaluation. However, as the practice has expanded, an unprecedented collection of legal issues has arisen along with it, making compliance on behalf of industry practitioners evermore nuanced.

Scope of this paper

Part 1 of this paper will outline distinct types of technology used in recruitment and contextualise the legal considerations associated with them.

Part 2 and 3 of this paper will then focus on the two most prominent areas of the law for consideration while employing automated decision-making software in recruitment- employment specific bias, and privacy and data protection.

Technologies in AI recruitment

People often think of automated decision making and conjure images of a rigid merit-based grading system, assigning numerical value to candidates.

However, in reality, the modern AI recruitment process comprises a series of interrelated technologies that fundamentally streamline existing recruitment processes in order to optimise them.

While identifying the technologies, it is important to understand the conceptual separation between the ‘pre-employment’, ‘post-offer’ and ‘post- employment’ stages of recruitment, and the stage in which each technology is applicable. This is as each stage, (and hence the technology) is held to slightly varying legal standards concerning both employment and privacy law.

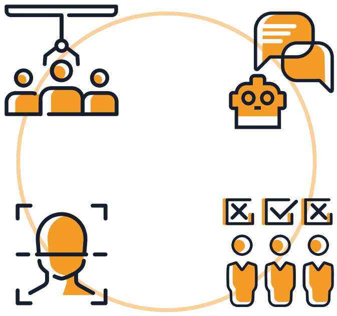

AI Matching Technologies

What it is:

Used in the preliminary sourcing stages of the recruitment process, they evaluate potential candidates based on relevant criteria, as defined through a role-based analysis. Purported as documented and data-driven, matching technologies are often hailed as a way to combat human biases.

Associated areas of legal consideration:

The impartiality of AI is often obstructed by the bias ingrained into the research with which they are developed. This presents a significant employment law consideration as it might essentially replicate or even exacerbate discriminatory hiring lines.

Being strictly in the pre-employment stage, matching technologies must also be carefully designed with privacy by design principles in order not to encroach on the private information that isn’t freely provided by a candidate.

Chatbots

What it is:

Used as part of both screening and evaluation stage processes, chatbots are an efficient way to discern the compatibility of a candidate. Using algorithmic and learning-based response technologies, the chatbot is supposed to act as an objective interviewer.

Chatbots can even exist as simple FAQ engines or other simple procedural software such as scheduling and follow-up discussions.

Associated areas of legal consideration:

Privacy watchdogs increasingly raise concerns as to the threshold of privacy rights for candidates, in regards to what can and cannot be asked at a pre-employment stage, in line with prevailing privacy and employment regulation standards.

AI Grading Software Tools

What it is:

Algorithmic grading software tools can be conceptually separated into their purpose at each stage of recruitment. E.g. candidate compatibility analysis, resume and other application support material analysis, candidate response analysis etc. While functionally similar to matching software, it is important to note their distinction as they bear significantly different legal considerations.

Associated areas of legal consideration:

In the pre-employment stage, there is a significant privacy concern as evaluation software of various forms might be able to discern private information that ordinary human analysis might not be able to.

Furthermore, in all stages there is also a significant risk of algorithmic bias in evaluation that must be regulated through human intervention, automated decision making software fundamentally lacking human empathy, social awareness and critical reasoning skills.

Facial Recognition and Voice Analysis Software:

What it is:

Biometric analysis software is commonly applied during video interviews for enhanced candidate evaluation. Facial scanning for micro-expressions and audio analysis such as tone/ language choice etc are used to analyse candidates and evaluate their compatibility with the role at hand.

Associated areas of legal consideration:

Such technology is a distinct legal consideration, as it seeks to enhance the traditional recruitment process by factoring in metrics that were previously unattainable. Existing employment laws are hence difficult to adapt as they were not developed accounting for it.

The very existence and application of these technologies have raised various privacy considerations and employment law considerations. The advanced analysis of biometrics is a significant privacy consideration on account of their status as highly protected data. Furthermore, the requirement for submission to this arguably invasive process (that might potentially also be tainted by discriminatory algorithms) is another complex employment law related concern.

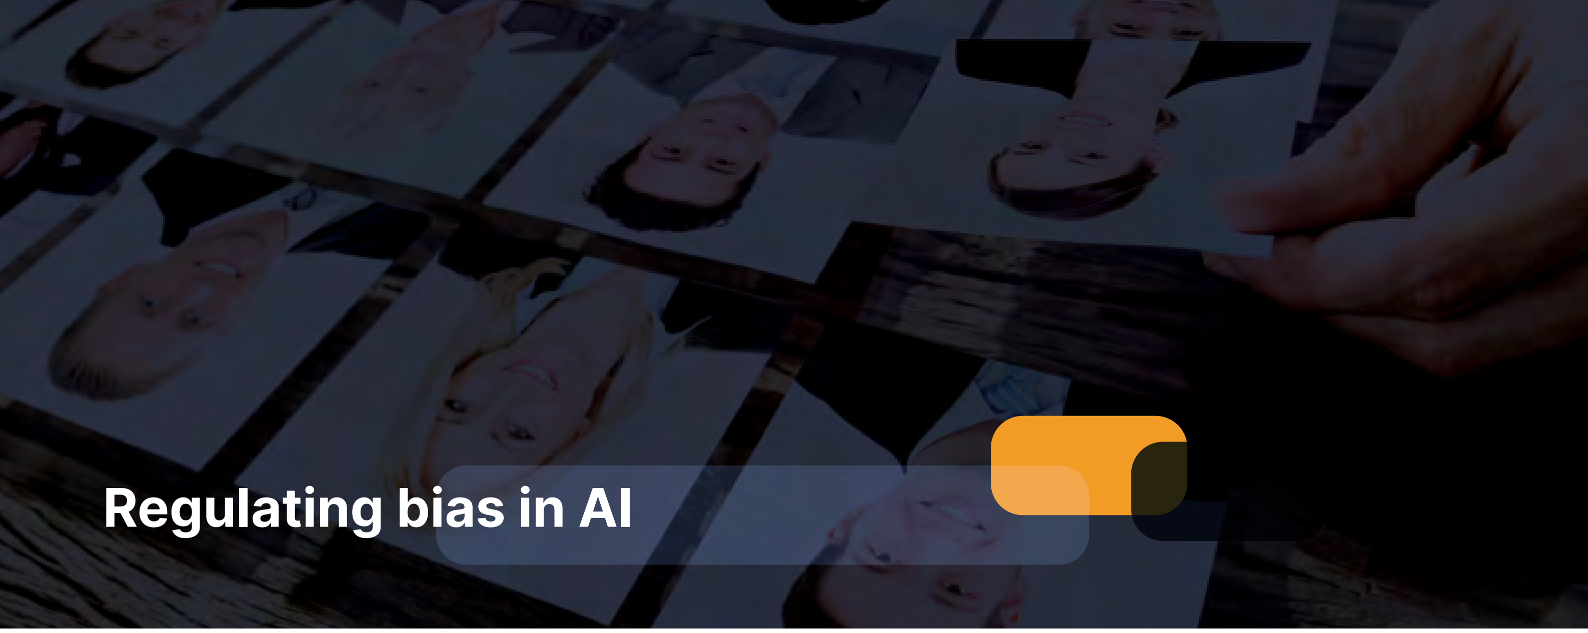

Regulating bias in AI

Overview of bias in AI recruitment

One of the fundamental concerns with AI decision making processes (not restricted to recruitment) is that the inherent impartiality of AI contingent on objective merit-based analysis is often obstructed by the bias ingrained into the data that they 'learn' from.

Legal discourse surrounding AI are imbued with the repetition of 'bias', 'discrimination' and unfairness'. Bias in traditional recruitment itself has been a prominent global issue through the last century, most countries instituting several laws to stem discriminatory hiring practices. The integration of AI and data-driven decision making into the recruitment process was initially hailed as a means to combat human biases, being objective, documented and merit-based. However, it has since been increasingly understood that AI's lack of contextualisation in decision making could potentially replicate or exacerbate these biases, raising concerns with the existing employment law standards.

The institution of regulatory efforts to mitigate AI bias in recruitment is quickly becoming evident in emerging jurisprudence, a 2020 Illinois state law being one of the first that requires explicit notice and prior consent from candidates being evaluated by AI. However, until more broadly applicable legal standards emerge, it is important to understand and adapt existing anti-discrimination employment laws in order to ensure legislative compliance and mitigate civil liability risk.

Types of AI biases

There are various distinct forms of bias in automated decision making. The primary being algorithmic AI bias or "data bias," At the root of this issue is the fundamental way that automated decision-making processes work – essentially by "learning" from large sets of data. Just like people, who's upbringing and experiences shape their biases, AI is subject to similar pitfalls. Rampant discriminatory practices are ingrained so deeply into our everyday socio-economic practices, that excluding these learnings from an AI's protocols has proven exceptionally difficult.

For several years, facial recognition software has boasted a high classification accuracy rate (over 90%). However, this fails to account for the fact that the facial mapping software tools used in the present day were primarily developed by Caucasian males, those subsequently being the first data sets that the machine learned, and the basis by which it continued to build its learning.

As the practitioners in the field have continued to diversify over the recent past, a growing body of new studies have demonstrated that this rating significantly erodes when concerning people of colour and gender, significantly a 34% higher error rate when mapping Black women. This being a demonstration of how the prevalence of bias in development can adversely impact practice.

However, tainted development models are not the only means through which automated decision-making systems can present bias. Unchecked data-driven algorithms can quickly replicate societal biases through evaluation of existing data, however much emphasis is placed on merit.

For example, Amazon Inc's machine-learning protocol demonstrated a significant issue, their new recruiting engine did not like women. Amazon's computer models were trained to evaluate applicants by identifying successful patterns in resumes submitted to the company over a 10-year period. Most came from men, a reflection of male dominance across the tech industry. Thereby, Amazon's system in effect "taught" itself that male candidates were preferable, penalising resumes that demonstrated any tangible connection to women's colleges or other such indicators of gender, as the algorithm understood this to be an indication of diminished compatibility.

While a jarring ethical consideration, the legal issue is less in the fact of the existence of bias, but rather in the effective regulation of it. Discriminatory intent and discriminatory effect each being distinct considerations in civil law.

Discriminatory intent vs discriminatory effect

While a potential civil liability nightmare, most legal regimes do not have the competencies to penalise the mere existence or use of potentially biased automated decision-making software in recruitment, as long as the said bias is accounted for and mitigated against. This is courtesy of the present legal paradigm that governs civil laws (that both the majority of privacy law and employment law fall within). Wherein, knowledge of potential harm without occurrence is not a sufficient basis for legal penalisation.

The practical benefits of the usage of AI in recruitment, have however rapidly outweighed the ethico-legal risks outlined above. Bias mitigation efforts such as third party bias assessments have become essentially customary in the application of automated decision-making systems.

Third-party evaluations themselves are an invaluable tool of legal indemnification. They allow employers to identify and mitigate any potential issues that might arise, thereby avoiding regulatory repercussions.

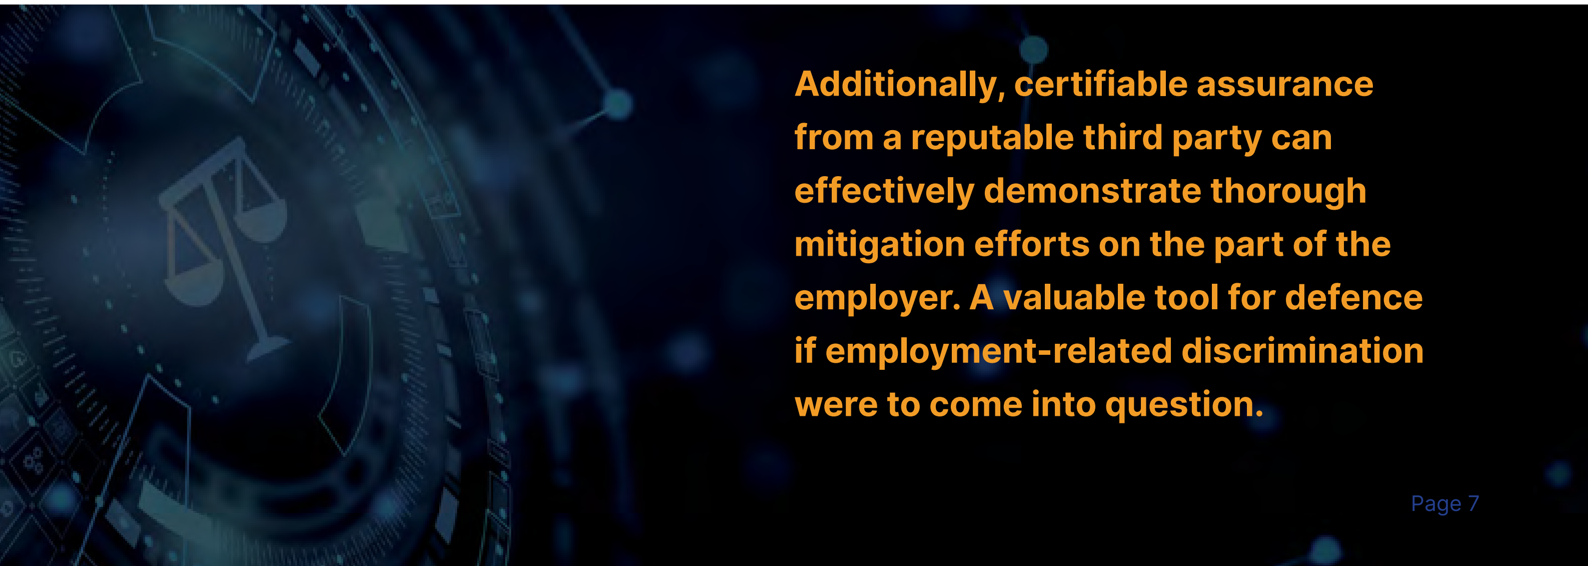

Additionally, certifiable assurance from a reputable third party can effectively demonstrate thorough mitigation efforts on the part of the employer. A valuable tool for defence if employment-related discrimination were to come into question.

Bias evaluations

The issue at the forefront in identifying a substantial legal standard necessary for a bias evaluation is that there is no concrete definition for what constitutes "fairness". Numerous mathematical criteria and models have been developed to define what it means for an algorithm to be fair. However, in a field with real-world applicability such as in recruitment, it is imperative that broader socio-economic contexts and exigences are accounted for.

This is often where the law factors into bias evaluations, as they are meant to reflect the societal norms of a jurisdiction (although that might not always be the case). The legal standard for bias combatting hence might be considered a suitable threshold to hold for 'non-discriminatory software'.

The EU, Canada, Australia, India and a majority of other states (particularly those under commonwealth influence), shield employment-related bias under the umbrella of basic human rights protections. Following the definition of non-discrimination in the UDHR with a generalised mandate of equal treatment to those of all races, colour, creed, gender and political or social opinion.

The US (and on a smaller scale the UK), are the only systems within which there exists an explicit statutory outline of a range of protected classes or characteristics. These classes being entitled to enhanced protections as a result of historic systematic oppression.

In the US, this umbrella of Title VII "affirmative action" protections essentially extends to every class excluding Caucasian American Heterosexual Males (who have long since dominated the majority of the socio-economic space).

The jurisdictional definitions of 'bias', 'fairness' or 'equality' continue to evolve and disaggregate rapidly. Ensuring up-to-date standards of these is an important compliance consideration for any AI system.

Regulation of bias evaluations are also rapidly emerging. In New York State, it has recently become a mandatory measure for all AI recruitment. Requiring that these bias evaluations must be conducted by a domestic third-party evaluator that holds accreditation within the jurisdiction. While a compliance necessity in New York, a bias evaluation is also valuable for general legal indemnity in numerous jurisdictions. It must be accounted that as algorithms continue to learn from data, and regulations change, bias evaluations must be conducted periodically in order to reflect the present best practices.

The jurisdictional definitions of 'bias', 'fairness' or 'equality' continue to evolve and disaggregate rapidly.

Privacy considerations for AI recruitment software

As data-driven algorithmic decision making has become increasingly prevalent, the primary concern around the collection, retention and processing of this data, has been the privacy rights of the individuals whose data is at hand. Employment law and privacy law have always gone hand in hand. Even the EU's General Data Protection Regulation (GDPR) is being drafted explicitly highlighting consumers and employees as distinct protected groups. Hence for a field intertwining both prospects as heavily as recruitment does, privacy and data protection are critical considerations in ensuring legal compliance.

The privacy notice

As the debate on privacy rights has developed its body of jurisprudence, notification and consent have become the essential compliance measures that have defined 'Privacy' and 'Data Protection' within customary international law. The customary mode for notification has since rapidly become a comprehensive privacy notice.

The introduction of the GDPR and the requirement for "comparable protections" as a prerequisite for data transfer has meant that a large part of global privacy jurisprudence is aligning with GDPR rules. This has essentially created a quasi-standardised privacy notice within comparable fields or technologies.

The key areas in which regulatory standards remain disaggregated and hence privacy notices become divergent are in the acceptance of implied consent upon notification, thereby categorising systems into 'opt-in' and 'opt-out' regimes. As per Article 22 of the GDPR, The EU requires explicit notice and opt-in consent for any automated decision-making processes. Certain states also require that a privacy notice make explicit reference to legal rights that a protected individual holds within the jurisdiction.

With the privacy notice being held on a pedestal by data controllers as a tool for legal indemnity, and with the rapid development of technology; such disclosures are becoming increasingly longer and more complex. The issue that subsequently arises is in the justification of informed consent, the basis for which is potentially jeopardised by information overload (infoxication) brought on by drawn-out privacy notices.

The premise of informed consent is, that by agreeing to utilise a service, an individual does so armed with the knowledge of why their data is being collected, and the means with which it will be retained, processed and shared. However, studies have demonstrated that less than 1% of internet users actually read privacy notices, fundamentally invalidating the premise of said consent.

Just in time notices

The legal standard for valid consent essentially holds that it must be informed and knowingly given. In essence, an individual must understand that they have consented to something and the basis on which they have provided this consent.

In the balance of this consideration, several legal systems have begun to move their focus away from the elongated privacy notice. Instead, calling for broader and more simplified notification throughout the data collection process.

The UK and Canadian privacy regimes, as an example, endorse a meaningful consent guideline. Following the idea that one-stop notification and consent are fundamentally redundant, they suggest notification in the form of just in time notices. This model is better suited to facilitating informed consent, through increased accessibility of information.

The use of just in time notices is also rapidly increasing in popularity as best practice for data controllers in several other regimes. Notably, in the US, where the threshold for civil liability is notoriously low, the insurance of adequate notification through just in time notices have become an invaluable asset for legal indemnity.

Additional safeguards for biometric data

Healthcare technology has long since been developed with privacy by design principles emphasizing special protections for biometric data, as reflected by long-standing legislation such as HIPAA. The introduction of the GDPR provided one of the first comprehensive and authoritative definitions for the scope of what constitutes biometric data in Recital 35.

More importantly, it highlights biometric data as being within a special category of protections. Thereby prohibiting the processing of biometric data, unless with the explicit consent of the data subject, or if necessary for public interest.

This is due to the higher risk threshold associated with the handling of biometric data. The misuse of an individual's facial features and fingerprints bears significantly more repercussions than the simple misuse of their email address or personal particulars.

Furthermore, in addition to data protection concerns, privacy concerns have been raised at the mere requirement of biometric analysis for employment purposes. The concern is rooted in the ability of the facial recognition software to discern additional details that users are not required to and do not consent to provide at a pre-employment stage.

For example, certain software could process data in enough detail to discern personal information regarding a candidate that would not be discernable in an average interview, such as age or sexual orientation. If this information is not voluntarily provided to the employer; this form of data processing could easily be considered a breach of privacy.

Even if this information was not factored into the evaluation of the candidate, this would be legally immaterial as if there was no consent to its initial collection, this would still be in contravention to the GDPR.

Conclusion

The introduction of AI to the recruitment process poses novel challenges for the law. Until comprehensive legislation is developed to govern AI itself, practitioners are left to apply traditional legal doctrines to complex and potentially unexplainable systems.

In relation to employment law, the primary concerns revolve around bias producing or exacerbating algorithms. A large part of this risk is rooted in misguided and unsupervised mandates, such as seen in the Amazon example.

Unsupervised pattern centric learning systems, such as in the above example, pose significant legal concerns when practised for recruitment purposes. Employment law compliance must prioritise merit-based analysis rather than pattern repetition.

At impress.ai, we employ a combination of rules-based and supervised learning algorithms. Where the rules are set based on Organizational Psychology research, this approach demonstrably combats bias as opposed to the earlier example, as it evaluates the merits of a candidate within a prescribed mandate, rather than replicating potentially problematic hiring norms.

Even with merit-based algorithms, it is imperative to routinely perform bias assessments by accredited bodies. As one would periodically perform employee evaluations, intelligent decision-making systems must also be assessed to ensure that their mandate remains impartial and in accordance with prevailing regulations.

Privacy considerations for AI recruitment, as for most data controllers, are imperative and aplenty. A significant part of the legal challenge in this area is rooted simply in the fact that these technologies present a novel consideration to the right to privacy itself. Advancements in data processing offer the means to collect data that would have been previously unattainable. It is hence imperative for data controllers to ensure all privacy-related notifications are

a. easily accessible,

b. comprehensible and,

c. up to date.

While this white paper sets out strenuous objectives, it is written in recognition of existing best practices, and the direction in which the regulatory standards applicable to AI may evolve in the coming years. While there is no legal requirement to institute all of these measures at present, proactivity might provide practitioners with retroactive indemnity and a competitive edge.

impress

Interested in more information?

Contact impress.ai

✉ contact@impress.ai

🌐 impress.ai

📍 Head Office, #08-01, 80 Robinson Road, Singapore- 068898

About impress.ai

impress.ai is a leading recruitment automation solution provider with a focus on making accurate hiring easier. Powered by AI, impress.ai's intelligent recruitment automation platform enables businesses to streamline their end-to-end recruitment process. impress.ai helps enterprises screen, engage, and hire the best talent with accuracy, consistency, & efficiency. We have partnered with leading businesses globally, offering 24/7 recruitment capability, helping them qualify the best candidates, increasing their hiring efficiency, and improving employee retention while consistently delivering superior candidate experience.

Headquartered in Singapore, impress.ai has a regional presence in the USA, Australia, India, and Indonesia. impress.ai was accredited by IMDA under the Accreditation@SG:D programme and has won 'Silver' in the Most Promising Innovation category at SG:D Techblazer Awards 2020.

impress