

Contents

Vendor claims for AI	3
What is AI?	4
What is an algorithm?	6
How is AI used in Human Resource Decisions?	7
Using AI in HR - the Research	9
What are the risks of AI in HR processes?	11
Garbage In garbage out	11
Findings lack validity	13
Predictor variables not validated	15
Findings are biased	16
Entrenching the status quo	18
Privacy and fairness Issues	19
Lack of Transparency	20
Positives and Provisos	21
Possibilities	21
Practicalities	21
Is AI (ML) any better than traditional methods and analysis?	22
A final question – is AI really helping us?	23
References	24

Right People in the right place at the right time?

Can AI provide the answer?

Vendor claims for AI

“AI- powered” systems are the next shiny object for start-ups, but also for established and major technology companies.

Nowadays almost every piece of software aimed at HR seems to have an AI aspect with all sorts of claims – from various websites;

“AI enables HR teams to extract insights from data and give recommendations in real-time”
‘AI will change the recruiter role through augmented intelligence which will allow recruiters to become more proactive in their hiring, help determine a candidate’s culture fit’
“AI in HRIS software can offer that completely unbiased perspective.”

But is this really true?



“AI in HR software can help a company’s HR team make better decisions. Artificial intelligence in HR can improve reporting, analysis, and even predictions for the future. Artificial intelligence in human resources management can be involved in important processes such as recruiting, onboarding, employee retention, performance, analysis and more”

What is AI?

AI in the context of Human Resources software is not really Artificial Intelligence – it is Machine Learning.

Machine Learning is a computerised method that looks for patterns in real world examples. The patterns are used to make predictions.

These predictions are specific to the example context.

The real-world examples are the historical data used to train the machine. Machine Learning requires a **LOT** of training data.

The methodology has its own terminology.



Random forest
Hyper-parameter tuning
Holdout set
Tree based model
Regularized regression
Feature selection

There are various methods for programming the machine to learn. The main ones are;

SUPERVISED MACHINE LEARNING

This process uses data sets to predict an outcome of interest when new data is input. For example, using applicant data to predict job performance for new applicants. The training data is existing applicant data and the job performance data of appointed applicants.

NATURAL LANGUAGE PROCESSING

A variation on Supervised Machine Learning where the machine learns to categorise, guided by humans, based on word frequency, position and semantic variations.

This process involves detecting patterns in complex information. It can be used to find previously undetected relationships. For example, the characteristics of those who are working remotely. This method just groups the cases in the data there is no target outcome prediction.

UNSUPERVISED MACHINE LEARNING

A type of machine learning in which the training data has no labels or scores. The machine discovers naturally occurring patterns in the training data set. It may cluster training cases data into categories with similar features and/or find features that are useful for

distinguishing between cases. There is no target outcome prediction.

Unlike humans, Machine Learning is not able to identify abstract principles and apply them to a situation that is different from the training data set.

What is an algorithm?



An algorithm is ‘an “unambiguous” specification in relation to problem solving’.

Unambiguous means that each step in the process is identified, inputs and outputs are defined, and the end point produces the correct result.

Algorithms are used as specifications for performing calculations, data processing, automated reasoning, automated decision-

making, and other tasks. For example, the process for doing long division.

Algorithms can be deterministic – that is represent a causal relationship. If A occurs, then B follows – every time.

They can be probabilistic – where the occurrence of A increases the probability of B occurring – but does not guarantee it.

Most HR problems are probabilistic. Traditional HR statistics use regression models which are also probabilistic.

Algorithms that seek to solve HR problems are therefore heuristics. Defined as an approach to problem solving that may not be fully reasoned out step wise. It may not guarantee correct or even optimal results, especially in problem domains where there is no well-defined correct or optimal result.

So – Is it safe to use AI for HR decisions?

How is AI used in Human Resource Decisions?

There are potentially a huge number of data sources that can be used to predict factors of interest to HR, for example job performance, employee turnover, retention, and job satisfaction.

By far the most frequent application is in the hiring process.



RECRUITMENT

The recruitment process can be split into 4 stages – sourcing, screening, interviewing, and selection. All these stages are being automated now as organizations try to respond quickly to large volumes of applications.

SOURCING

Theoretically AI can be used to remove bias in job wording to encourage a more diverse group of applicants

APPLICANT SCREENING

Many organizations have turned to AI to help screen job applicants. It is likely that over three quarters of resumes sent to US companies are not read by humans at all. In theory ML is better at screening resumes than humans for key words linked to job requirements. It is certainly faster.

INTERVIEWING

AI is being used for interviewing and video interviewing. AI can automate structured behavioural interviews.



It is used for video interviews, where candidates are asked to record answers to particular questions and to provide open responses highlighting their strengths. These videos are then algorithmically analysed, not just for words but for voice tone, facial expressions, and body language.

SELECTION

AI can be used for analysing answers to questions to identify candidates with characteristics that top performers have.

These question-based assessments can include personality tests, situational judgment tests, and other formats. This is like the process used in psychometrics and assessment centres - but without independent validation.

OTHER USES

Some vendors promote the use of AI for career pathing and training recommendations, based on skill matching.

ML is also used to investigate predictors of job satisfaction, engagement, and staff retention.

Does AI really work – where is the evidence?

Using AI in HR - the Research

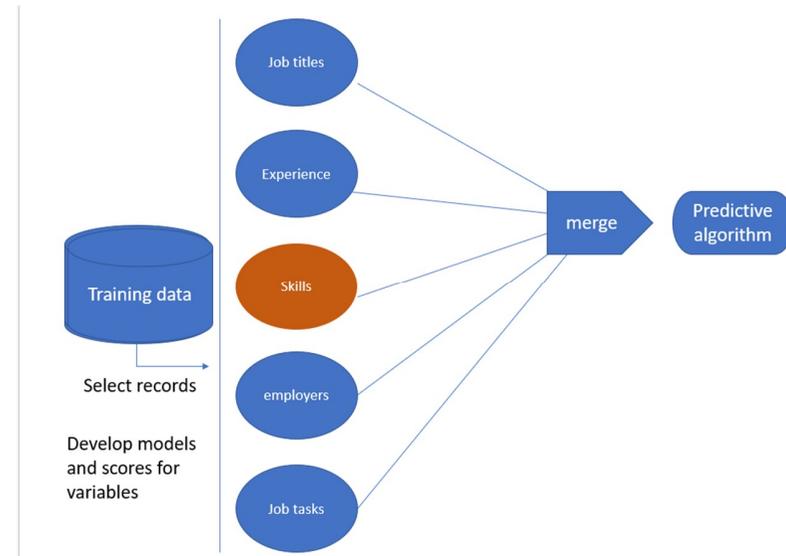
In fact, there is almost no independent research into the effectiveness of AI/Machine Learning in supporting Human Resource decisions.

A study using a large US government agency investigated the effectiveness of using natural language processing (NLP) to rate applicant data¹. This was a task already performed by expert human raters.

Humans use a wide variety of words to describe the same concept. They often use the same word to refer to different concepts in different contexts. This is of course a key issue in scanning resumes.

NLP looks for word frequencies, positioning, and semantic similarities. Even if target key words like 'leadership' are not used in the resume, leadership skill can be inferred from

the frequency of similar terms such as 'initiative' and 'organizing'.



Step 1 - train the machine - create models from existing data

The machine trained on over 40K historical applications with extensive human guidance.

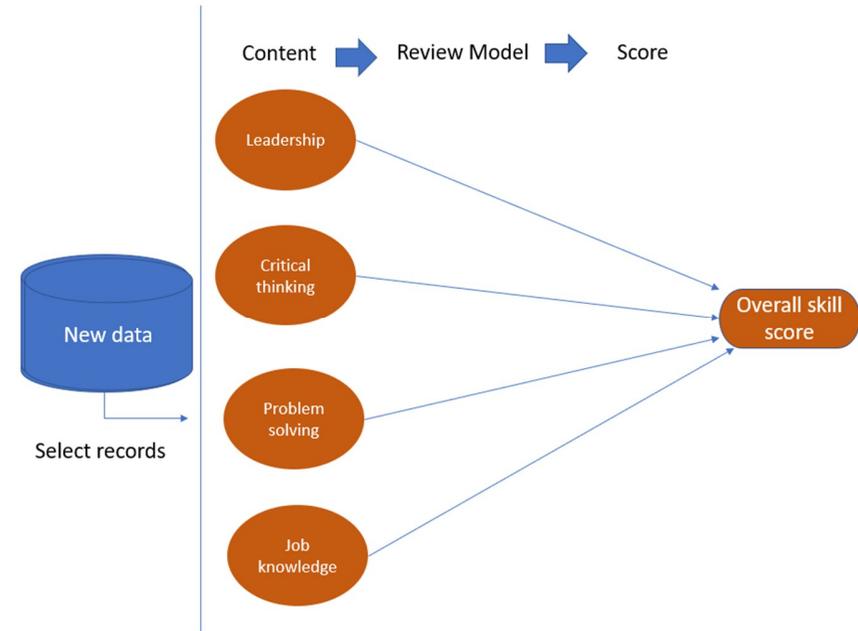
The factors assessed were 6 core competencies used to assess if candidates should be

interviewed. The data analysed included a 200 word “Achievement Record” essay, education, work experience, other skills and background on the application form and scores on a test battery.

For each competency between 1200 and 1440 categories were extracted by the ML. 7-900 from the essay, 60-90 for job titles, work duties 1-300, employers 50-130, special skills 60-90 and other experience 10-20.

At the end of this the AI equaled the human decision making with the error rate not significantly greater than the 3 human raters normally employed.

A 2019 study investigated the use of machine learning to predict job performance from work history². The 3 factors examined were job related knowledge and skills (derived from matching job titles in historical applications to Onet jobs and skills listings), length of tenure and reasons for leaving. The training was extensively human guided. The results showed that all three factors were associated with job performance in the expected directions, but the correlations were very weak.



Step 2 - Use the trained models to analyse new data

What are the risks of AI in HR processes?

Most of the algorithms used in the HR function are probabilistic or descriptive, rather than based on causal relationships tested in the real world. Therefore, relationships discovered may not be suitable for decisions that affect peoples' employment.

Garbage In garbage out

Machine learning depends on training data. If this data is incomplete and or incorrect then the models/algorithms derived from it will not produce useful predictions or relationships.



The training data should include all the factors that predict good performance. The data format must be categories or numbers. The new data input must match.

Of course, the only training data available is historical. It must be from the same context and a LOT of data is needed. Realistically this means only large organisations will have sufficient data. Others may have to use vendor pooled data – and since this comes from different contexts the accuracy of the results will reduce.

The data may not represent the full spectrum of candidates – for example more males than females apply for tech jobs.

The target outcome must be objectively measurable;

- a definition of the ideal applicant based on a concept put together from indicators measured after hiring – usually job performance
- A target number – for example sales achieved
- Yes/No – whether or not an applicant with these characteristics has been previously hired, usually based on interviews, psychometrics etc.

The problem is evaluations of job performance are highly subjective. Few jobs have numerical performance achievements.

This performance data is limited- there is no data on how those who weren't hired would have performed.

In practice, it is impossible to even define, let alone collect data on, **an objective measure** of a “good” employee.

So the more subjective and unreliable the measurements of the ideal hire – the less accurate the machine prediction.

Findings lack validity

According to Arvind Narayanan Associate Professor of Computer Science, Princeton ³.



“much of what’s being sold as “AI” today is snake oil – it does not and cannot work”

AI does pattern recognition well – so it is great for face recognition, diagnostics from scans and

x-rays, speech to text, matching music and images, even producing “DeepFakes”. Here the problem is with the high accuracy creating ethical concerns.

Its use in the tech sector for judgements on what is spam, hate speech, copied material, for recommending content on search is a problem too.

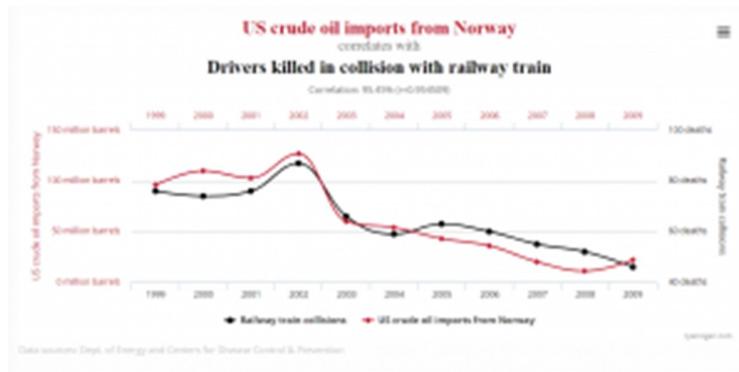
So is use for automatic grading of student essays. (See above the low predictability from analysis of applicant essays). Whilst accuracy is increasing, it is nowhere near perfect and that also leads to ethical concerns.

Narayanan highlights applications of AI to make social predictions as those that are really dubious. Their inaccuracy and use in high stakes situations such as policing, social welfare and employment raises serious ethical concerns.

Predicting job performance and identifying staff at risk of leaving are squarely in this category.

Machine Learning brings with it the risk of spurious correlations in these huge data sets.

The correlations may be just co-incidental.



or

- (a) They do not reflect the true relationship – there are other factors that underlie the discovered relationship. For example, a person who is always on time to work fits the criteria for a top performer.

What may be overlooked is the ease of their commute and their family status – those with children are more likely to be disrupted.

- (b) Taken as causal in a particular direction when the opposite may be true.

For example many different articles suggest that employee satisfaction is one of the strongest predictors of long-term positive company performance.

However, the relationship could well be reversed – in that people get more job satisfaction when they perceive they are working for a successful company.

Findings are biased

The results of a ML algorithm should be independent of variables that are discriminatory – such as gender, age, sexuality, race, religion, politics or ‘proxy’ variables for any of those. (Sport for example is a proxy variable for gender because more males are involved in sporting activities.)

Bias can be ‘inherited’ from training data – hiring decisions made in the past.



Further, as was the case with Amazon, the training data may be imbalanced – in that case, as in many IT companies, the training data reflected a workforce with more males, and created a bias against females.

“Bias is caused by data selection, the right to opt out of data collection, existing bias in decision making and inappropriate choice of algorithm.”

Professor David Parry, Auckland University of Technology

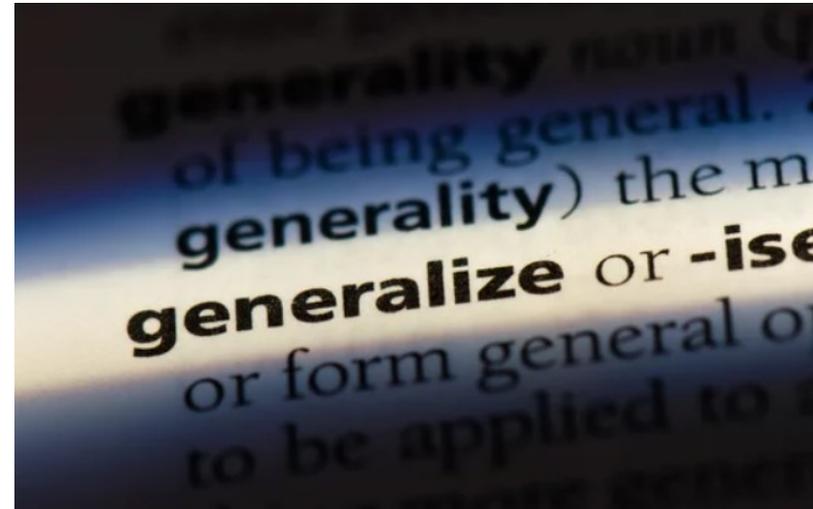
In a study of 18 vendors offering AI supported hiring ⁴ it was found that all but 3 claimed measures to prevent bias or adverse impact. None explained how.

Eight build assessments based on data from the client’s past and current employees

Clients determine what outcomes they want to predict, including, for example, performance review ratings, sales numbers, and retention time.

Eight do assessments “somewhat” tailored for clients but not using client data. The assessment scores applicants on various competencies, which are then combined into a “fit” score, based on either vendor or client input, to create a custom formula.

Three do ‘pre-built’ assessments of commonly used factors. These provide assessments for a particular job role (e.g., salesperson), or provide a sort of “competency report” with scores on a fixed number of cognitive or behavioral traits (e.g., leadership, critical thinking, teamwork). that have been validated across data from a variety of clients.



Experts analyze the job description and role for a particular client and determine which competencies are most important for the client’s needs.

Pre-built assessments must by nature be general, but consequently, they may not adapt well to the client’s requirements.

In practice, there is little difference from traditional selection tools.

Entrenching the status quo

A major criticism is that the use of AI risks encoding the status quo, bias included.



The need to use historical and current data for training the machine assumes that the organization's context will carry on into the future.

It cannot allow for changes in circumstances from mergers, acquisitions, changes in market environment.

There is no scope for improvement, the machine doesn't know what it doesn't know – there is no data about the kinds of people that were not hired, that may have been successful.



Privacy and Fairness Issues

A key part of 'privacy' is control over one's own personal data.

Traditionally, applicants have been able to choose what information about themselves to provide a potential employer to present themselves in the best light.



The scraping of data from social media, emails, resumes and other internet sources as well as the pooling of personal data for use across client organizations posed problems, especially in the EU.

The use of machine learning in video interviewing applications to interpret facial expressions, voice tone and body language has been called into question. One issue is inaccuracy – for example facial recognition is less accurate for African Americans than for white people, and less for females than males. Another issue is lack of specific consent for how it will be used.

Some US states are now enacting legislation to require informed consent before an employer bases a selection decision on data beyond the applicant's control.

AI enhanced selection methods are dynamic. The normative data set continuously changing as new applicants enter the pool. The problem this causes is that applicants are evaluated on different variables depending on when they applied. This 'disparate treatment', makes organisations vulnerable to legal challenges.

Lack of Transparency

The major problem is that most AI-based selection tests are not transparent as to how the tools were developed.

Nor is it clear how the algorithms score applicants against the target criteria.

EU privacy regulations – the GDPR, require that any models or algorithms used for decision-making must be explainable. A stance increasingly adopted by other jurisdictions.



There are three aspects of transparency required in employment decision making to defend against claims of adverse or disparate treatment;

- It must be clear what variables are used e.g. which aspects of resume, personality tests, work samples
- It must be clear how applicants are rated
- It must be explainable – feedback can be given on why a candidate was not taken forward

so a model that humans can understand.

As we have seen many AI HR solutions don't meet this requirement.

Positives and Provisos

Possibilities

In summary the promise of AI in Human Resources is to complete some operations, especially in the hiring process, more quickly and efficiently than humans can. This gives organizations a better chance of attracting, identifying and hiring the best people in a competitive labour market.

The use of big data may uncover hidden relationships – for example the over specification of qualification requirements – that people without advanced degrees perform just as well or better than those with them.

Practicalities

There are very significant risks in using Machine Learning technology to make predictions of likely job performance. Key risks are inherited bias and lack of fairness,

entrenching the status quo, lack of validity of predictor variables, privacy issues and lack of transparency for HR staff and applicants.

Is it worth it?

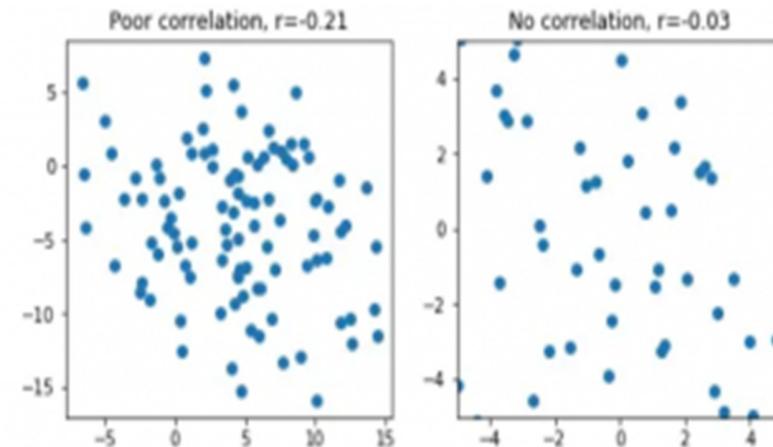
Is AI (ML) any better than traditional methods and analysis?

In contrast to classical statistical analysis which uses just a few variables, Machine Learning algorithms can handle huge numbers of predictor variables. For example, the capture of mouse click and cursor movements in a game.

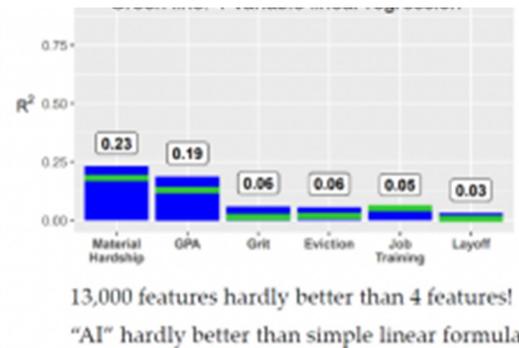
Despite this, studies comparing the output of traditional statistical analysis with machine learning have found little difference. A recent study⁵ used data on children from birth to 9 years gathered from interviews with over 4000 families. Information on almost 13000 factors was analysed to see what might predict 6 outcomes at age 15.

- School grade point average
- Personality factor of “Grit”
- Household eviction
- Household material hardship
- Primary caregiver layoff
- Primary caregiver participation in job training.

Prior to the age 15 interviews, some 160 research teams used various ML techniques in a competition to see whose prediction could be the most accurate. Overall, the best ML predictions were not accurate when the results for the age 15 interviews were in. For child grade point average and household material hardship a very weak correlation of around 20% and for the other factors correlation only of about 5% – just a random relationship – as the scatter plots show.



Worse analysis by traditional statistical methods got almost the same answers – as Aravind Narayanan sums up in his presentation “How to Recognize AI snake oil”



A final question – is AI really helping us?

Some ⁶ are asking whether we are allowing the use of AI to devalue human judgment – the use of “personal experience, tacit knowledge, intuition, creativity, insight, empathy, trust”, and so on.

These authors suggest that AI/ML systems overestimate the world’s predictability and underestimate the novelty of situations.

“Unchecked, the spread of these complex tools both introduces new forms of fragility and instability, as we saw with the 2008 financial crisis, and works to deprive us of the skills and know-how to intervene.”

References

1. Campion et al 2016. Initial Investigation Into Computer Scoring of Candidate Essays for Personnel Selection. *Journal of Applied Psychology* 2016, Vol. 101, No. 7, 958–975 Michael C. Campion Michael A. Campion Emily D. Campion Matthew H. Reider
2. Sajjadiani et al 2019. Using Machine Learning to Translate Applicant Work History Into Predictors of Performance and Turnover. *Journal of Applied Psychology* 2019, Vol. 104, No. 10, 1207–1225 Sima Sajjadiani Aaron J. Sojourner, John D. Kammeyer-Mueller, Elton Mykerezi.
3. Aravind Narayanan November 2019 – Presentation How to recognise AI Snake Oil. Centre for IT Policy Princeton University.
4. Raghavan, M., Barocas, S., Kleinberg, J., and Levy, K. (2020). Mitigating bias in algorithmic hiring: evaluating claims and practices. In *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 469–481). Association for Computing Machinery. <https://doi.org/10.1145/3351095.3372828>
5. Matthew J. Salganik et al 2020. Measuring the predictability of life outcomes with a scientific mass collaboration <https://www.pnas.org/content/pnas/117/15/8398.full.pdf>
6. Joseph E. Davis, Paul Scherz 2019. Persons without Qualities: Algorithms, AI, and the Reshaping of Ourselves. *Social Research: An International Quarterly*, Volume 86, Number 4, Winter 2019, pp. xxxiii-xxxix (Article)