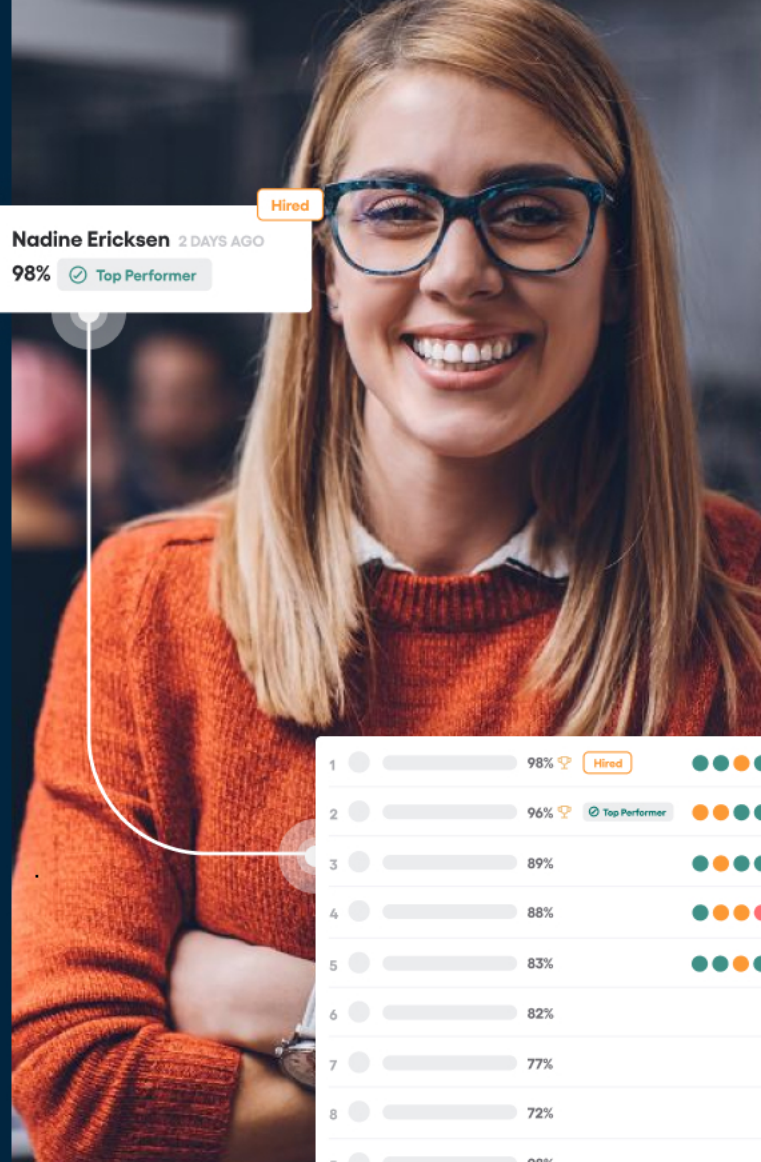# vervoe

Understanding our AI

At Vervoe we believe artificial Intelligence (AI) in recruitment can be a double-edged sword. Used ethically, AI is a tool for good that provides powerful, bias-free data points. But in an unregulated industry, some vendors are releasing ill-considered AI solutions that automate and compound human biases.

As AI technology becomes increasingly prevalent in recruitment processes, we think it's important that the people who use the solution understand the technology, how it works, and recognize the ethical challenges involved

Nadine Ericksen 2 DAYS AGO
98% ✓ Top Performer

1	98% 🏆 Hired		
2	96% 🏆 ✓ Top Performer		
3	89%		
4	88%		
5	83%		
6	82%		
7	77%		
8	72%		
5	98%		
6	96%		

AI in recruitment falls into two subsets:

Rule-based AI (Decision trees)

Rule-based AI involves very simple, rules-based engines that have been programmed by a human to respond in a specific way to a trigger. In hiring, an example could be software that scans CVs for keywords to create a shortlist. These decision trees, however, can only be as socially aware as their designers and there's been many examples of rule-based AI indirectly perpetuating inequality.

AI machine learning

Significantly more advanced than decision trees, machine learning (ML) learns from data, compares it to known outcomes, and finds patterns to understand the things that happened to reach a result.

Examples in hiring include ML models that learn from historical data what makes a good or bad hire in an organization, or AI that analyses candidates' written input and tone of voice.

The main difference is that rule-based AI works by doing what a human has told it to do, while machine learning figures out what to do for itself.

This is why at at Vervoe we use AI machine learning to predict the job performance of candidates, as its ability to continuously improve ensures equity and fairness for all.

How does AI machine learning work at Vervoe

We use a set of 3 different models that measure and predict performance.

How Model

Tracks and analyses the way a candidate interacts with the assessment.

For example: How long does it take them to complete? Do they scroll back to a previous question? Do they consistently type or are there long pauses? Do they click away and come back?

This model essentially grades how a candidate works and would be the same as you standing behind the candidates' watching them complete the assessment.

This model is the first of the 3 to provide a score and stops being dominant once the What model is activated.

What Model

Analyzes the content of the candidates' responses and benchmarks them against millions of other responses. Through natural language processing candidates' responses are viewed looking for words, phrases and sentiments that accurately reflect the outcomes required.

For example, if you're hiring a Call Centre Agent you might be looking for someone with skills like attention to detail and empathy. Our models process thousands of responses quickly, looking for certain words or sentiments that reflect these values accurately.

Candidates that display more of these are graded higher. The what model uses existing blind data sets, new candidate data, and user input in the form of correct answer samples.

Preference Model

This model uses the Naive Beyes method to predict probability. This model requires input from the user to train it to understand what the scale of bad to good answers looks like for their specific use case. Using a model called `iterative` to execute this, a user blindly grades a set of candidate responses giving a score from 0 - 10. The set of questions that are exposed for the user to grade are the furthest apart from each other.

For example, if a user grades one response as a 10, our model will then look for an answer that seems to be completely different to this to see how the user scores that answer. This variation in responses helps the model quickly identify and plug the gaps in between the potential score ranges to accurately grade all candidates with your preferences in mind.

How we combat bias

Machine learning can be susceptible to human hiring bias being introduced through the data set. If machine learning learns from biased signals like knowing whether an applicant is male or female, then the outcomes will have bias.
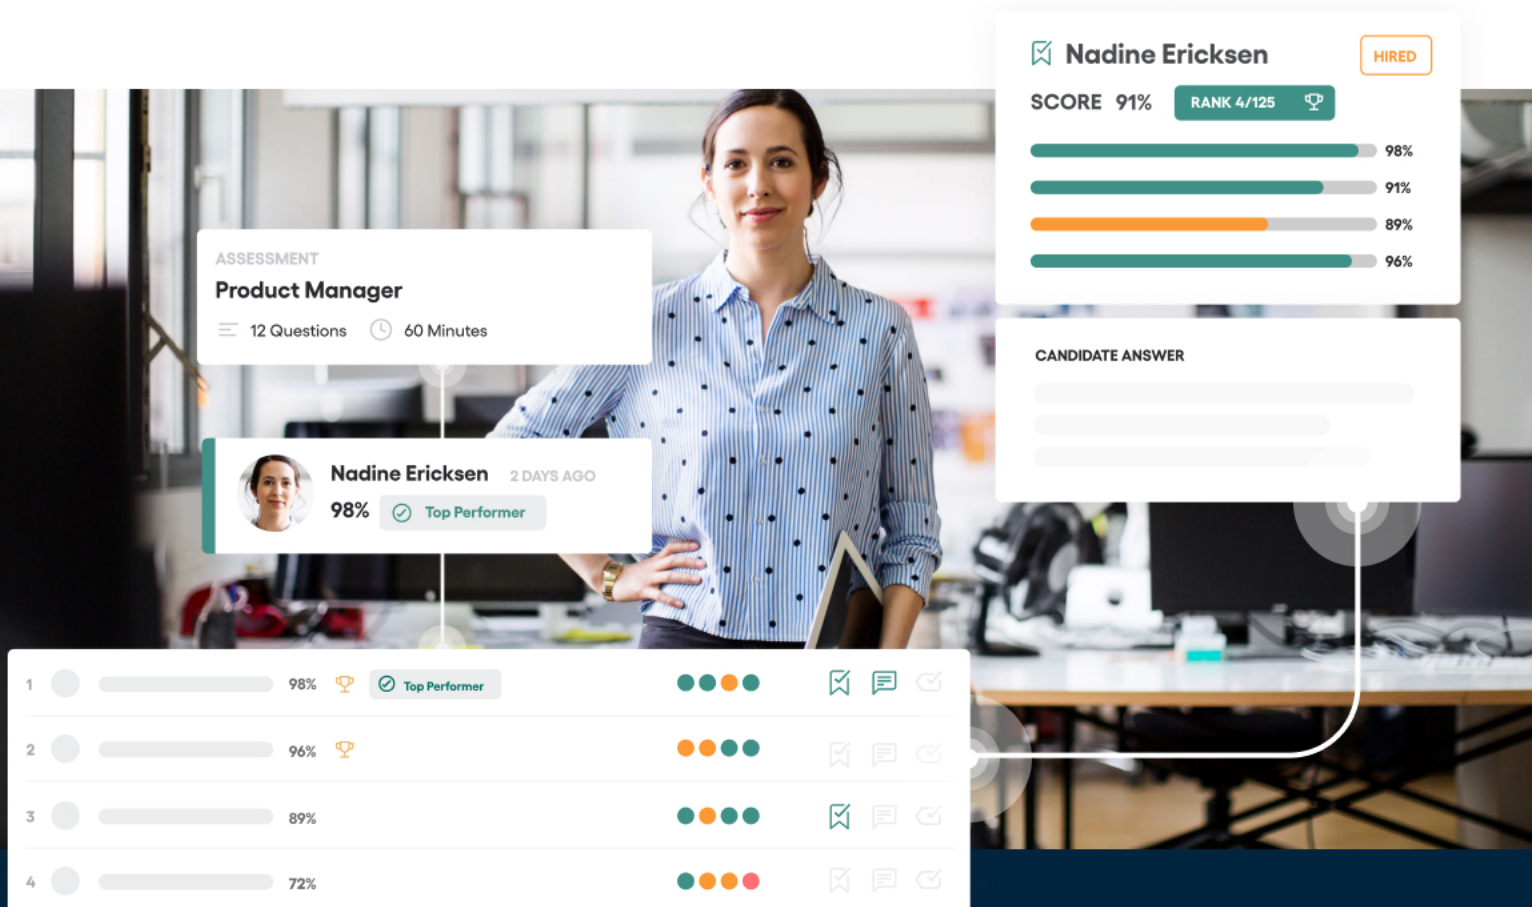
To prevent this occurring at Vervoe, we started with a clean dataset that we built up over a two year period with real candidates' applying for real jobs being graded by real employers.

Our models are also blind to anything that could potentially create bias. We've removed personal identifying information about the candidate from the assessment process.

We don't collect information like gender, location, ethnicity, background or education. Candidates are purely graded on the quality of their response data.

We also have checks in place that monitor manual grading by our users to ensure data integrity. For example, if a user is manually grading responses from their candidates but shows a clear bias towards anyone that mentions their children, we can detect that the drop in score is related to the keywords `kids` or `children`and flag this with the organization to investigate whilst removing it from our learning set.

At Vervoe we believe getting machine learning right from an ethical perspective starts with being human. We understand it's a huge responsibility to develop a system that helps to decide if someone gets hired and we treat that seriously. Our approach is to question what the AI is judging candidates on. Would we want to be judged on these things? If not then we don't include it in our data.

Nadine Ericksen · HIRED
SCORE 91% · RANK 4/125
98%
91%
89%
96%

CANDIDATE ANSWER

ASSESSMENT
Product Manager
12 Questions · 60 Minutes

Nadine Ericksen · 2 DAYS AGO
98% · ✓ Top Performer

1		98%	🏆	✓ Top Performer			
2		96%	🏆				
3		89%					
4		72%					

vervoe

Start testing candidates with Vervoe from day one

No lengthy implementation or training processes are involved in Vervoe set up.