Beamery

2023

35 Questions You Need To Ask Your HR Technology Vendors About AI

beamery.com

Private & Confidential - Do Not Share © Beamery Inc. All rights reserved

Introduction

A lot of HR teams are using AI these days... but some departments are concerned about the complexity of this technology, and the potential introduction of bias to their talent management processes. How can you be sure AI delivers real value, and solves recruitment and retention challenges more effectively and efficiently than another approach?

If you are at the start of your AI journey, or reviewing the solutions you have in place, here are the questions we would suggest you ask to ensure you are getting the most from any AI applications.



The Top 5 Questions

1. What are you doing to mitigate bias in your AI?

AI is trained using data sets, and if there is bias in the data, the outcomes from the AI will be biased too. Your vendor should be able to point to actions they have taken to reduce the chance of recommendations being biased.

2. Do you have an AI explainability document?

One of the fears people have around AI is that it is hard to understand: that decisions or recommendations could be made without anyone understanding why. Related to the reduction of bias, if an AI-powered solution is not explainable, there is a real chance of outcomes being unfair. A document explaining how the AI works is a great start.

3. What third parties have you been working with to audit your models and your training data?

With new regulations coming out of NYC, companies that use "automated employment decision tools" have to comply with some new requirements, including a yearly bias audit; and a summary of the results of such an audit made publicly available. Does your supplier offer this? Who is auditing their AI and what are the results?

4. How do you always ensure a human understands when they are interacting with AI?

As part of the new regulations mentioned above, companies that use "automated employment decision tools" also have to provide candidates with notice of automated processing; and allow candidates or employees to request an alternative process or accommodation. Does your vendor provide a way for people to consent to AI being used, and manage their preferences?

5. What are your AI ethical principles?

AI works best with human involvement: it is not there to replace HR people, but to save time on manual processes, and enhance the work they do. Your vendor should have principles in place to ensure that AI is used for good, across the talent lifecycle.

Understanding Data, Models & Experiences

Are you ready to dive a bit deeper? Here is the more detailed approach we would recommend when assessing an AI vendor in the HR space. We have divided these into the three layers of any AI technology: data, models, and experiences.

For each layer – data, models, and experiences – there are questions relating to Suitability, Objectivity, Explainability, Validity and Security.

DATA

AI starts with data. When it comes to the choice of data and data sets being used to train and power the AI model, it's important to drill down into what the intentions are. Is the vendor applying AI for AI's sake, or trying to solve a business problem?

You need the data to be useful, accurate, relevant and unbiased. While completeness is generally a good thing, you don't want to include more data than you need, and run the risk of compliance, security, or bias issues. At Beamery, our view is that the fairest approach to hiring (and development) is one based on skills. We use skills data to match candidates to roles, suggest employees for opportunities, or help you visualize career paths.



Questions About Data

SUITABILITY

6. **Why are the chosen data points used to train the models? Is all of it relevant for the use case?**

7. **What sources of data are used? Why?**

OBJECTIVITY

8. **How is the data checked for bias before it is used for training models?**

9. **Are there data points that may potentially cause bias, such as gender, race and ethnicity? How are these handled and why are these in the model?**

10. **Is the data representative of the distribution in the real world?**

11. **What is your approach to skills taxonomy, and how are you bringing together different data sources to create a uniform language around skills data?**

EXPLAINABILITY

12. **How is the data segmented e.g. by geography, industry?**

13. **How is the data stored and connected? Is it in tables, documents or graphs?**

VALIDITY

14. **How do you maintain data quality? Is it complete, fresh and deduplicated?**

SECURITY

15. **What consent was obtained when the data was used or acquired?**

16. **How is data per customer stored and used between other customers?**

17. **What steps are taken to ensure the data is securely stored and cannot be amended or stolen? Have you assessed for increased risk on data privacy breaches, on account of using AI (or Machine Learning)?**

In terms of storage, we use a Talent Knowledge Graph to explicitly model data points and build AI on top of that. This means that if we need to understand why a model is behaving in a certain way, we can look at how data is related, and see which data points are influencing the models. In other words, we can quite easily explain a recommendation.

Models

An AI model is a program that has been trained on a set of data (the training set) to spot certain types of patterns... and ultimately "learn" from the data, and apply that learning to achieve pre-defined objectives.

So, once again, asking about the models used helps you understand if AI is the right solution for the task at hand. You want to find out whether outcomes from the AI are going to be simple and explainable, how the models have been tested from bias (and how unfair biases have been mitigated or eliminated), and that they won't bring negative unintended consequences to your hiring process.

Ideally the vendor would be able to draw you a diagram of data flows: how the models connect, what are they trying to achieve, and the metrics they are using. How does it all work together to deliver on the desired outcome?



Questions About Models

SUITABILITY

18. Why were these models chosen over other models? What tests were done to choose these particular models?

OBJECTIVITY

19. How are the models checked for bias? Are there humans in the loop? How often are they checked for bias?

20. When there's feedback about biases or errors with the recommendations, how will this be taken into account in the model? What procedures and proofs do you have in place to ensure that this has been rectified?

EXPLAINABILITY

21. How does each model work individually and together to deliver on the use case?

VALIDITY

22. What quantitative and qualitative methods do you use to check if the models are working optimally?

23. Do you optimize for recall or precision? What about accuracy? Why? (If the solution is suggesting candidates for a role, are we aiming for the longest possible list, or the highest possible quality?)

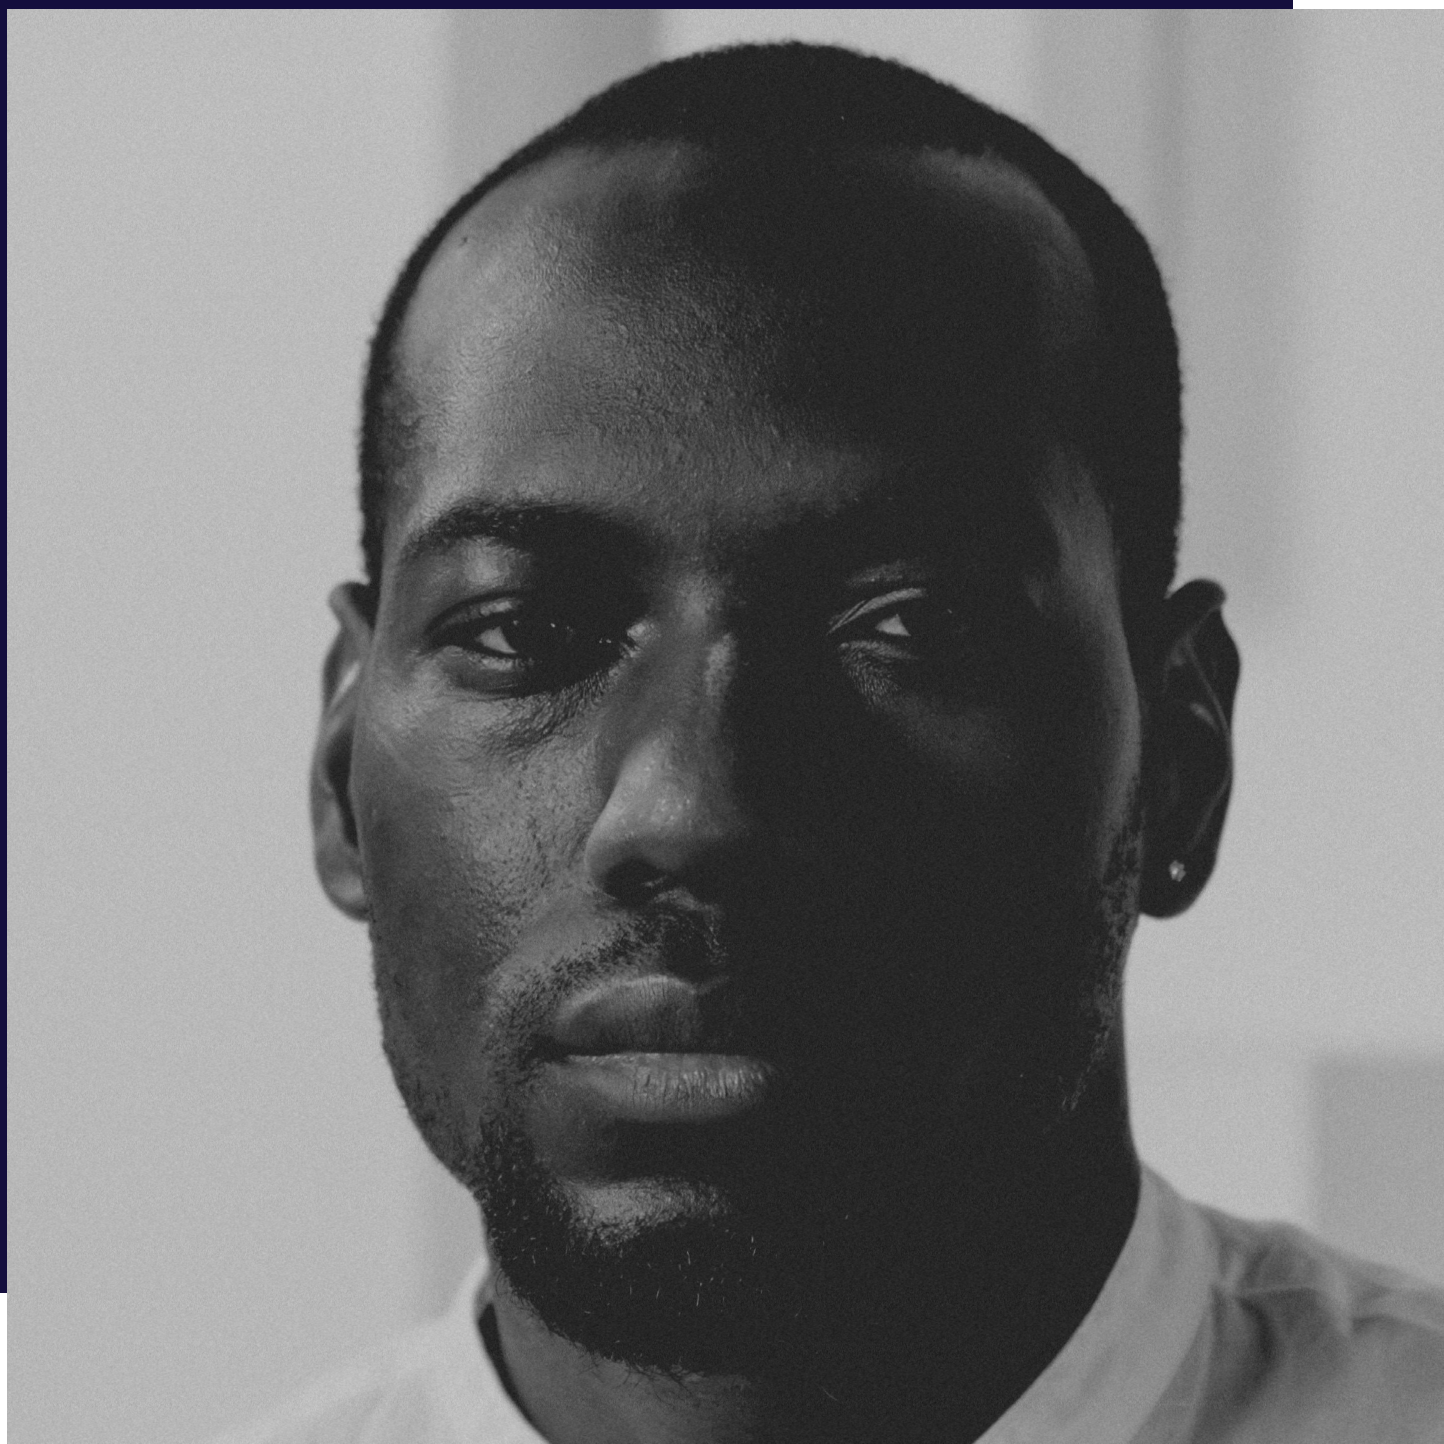
SECURITY

24. What steps are taken to ensure the models cannot be interfered with?

Experiences

When it comes to the experience for users (HR teams) and candidates (internal or external), it's the same question again: are we using AI for AI's sake, or is it actually adding value? You want to ensure they will have a better experience because AI is used.

While you don't want user interference to affect the model or add bias, we do need to get feedback from users on whether the right 'decisions' are being made, as well as being able to handle consent and preferences from candidates. And if someone is being recommended for a job over someone else, you need to know you can explain that 'decision'.



Questions About Experiences

SUITABILITY

25. How does the AI improve the user experience?

26. Where does it add value, and how does it enable users to achieve their goals more effectively than without the AI?

OBJECTIVITY

27. How much influence does the user have over the model?

28. How does user interaction and feedback on the models impact the model training?

EXPLAINABILITY

29. How is the AI explained to users? Is it clear on why the AI made recommendations, and the key factors driving recommendations? Is it easily understandable and reassuring?

VALIDITY

30. How is the AI tested as part of the user experience?

31. How do you measure success?

SECURITY

32. What steps are taken to ensure the AI experience is secure and cannot be interfered with?

	Suitability	Objectivity	Explainability	Validity	Security
Data	Why are the chosen data points used to train the models? Is all of it relevant for the use case? What sources of data are used? Why?	How is the data checked for bias before it is used for training models? Are there data points that may potentially cause bias, such as gender, race and ethnicity? How are these handled and why are these in the model? Is the data representative of the distribution in the real world? What is your approach to skills taxonomy, and how are you bringing together different data sources to create a uniform language around skills data?	How is the data segmented e.g. by geography, industry? How is the data stored and connected? Is it in tables, documents or graphs?	How do you maintain data quality? Is it complete, fresh and deduplicated?	What consent was obtained when the data was used or acquired? How is data per customer stored and used between other customers? What steps are taken to ensure the data is securely stored and cannot be amended or stolen? Have you assessed for increased risk on data privacy breaches, on account of using AI (or Machine Learning)?
Models	Why were these models chosen over other models? What tests were done to choose these particular models?	How are the models checked for bias? Are there humans in the loop? How often are they checked for bias? When there's feedback about biases or errors with the recommendations, how will this be taken into account in the model? What procedures and proofs do you have in place to ensure that this has been rectified?	How does each model work individually and together to deliver on the use case?	What quantitative and qualitative methods do you use to check if the models are working optimally? Do you optimize for recall or precision? What about accuracy? Why? (If the solution is suggesting candidates for a role, are we aiming for the longest possible list, or the highest possible quality?)	What steps are taken to ensure the models cannot be interfered with?
Experiences	How does the AI improve the user experience? Where does it add value, and how does it enable users to achieve their goals more effectively than without the AI?	How much influence does the user have over the model? How does user interaction and feedback on the models impact the model training?	How is the AI explained to users? Is it clear on why the AI made recommendations, and the key factors driving recommendations? Is it easily understandable and reassuring?	How is the AI tested as part of the user experience? How do you measure success?	What steps are taken to ensure the AI experience is secure and cannot be interfered with?

Additional questions

When it comes to assessing the quality of a new AI-powered solution, it is worth also considering Implementation, Adoption and Satisfaction.

33. **How is your solution actually implemented? Do you have any use cases? Will the experience of implementation be quick and delightful?**

34. **Is there a low barrier to usage? Is the AI intuitive enough to be a natural part of the workflow?**

35. **How happy are your current customers? Is the AI not only working but delighting users?**

Questions not to ask AI partners/vendors:

There are some questions that, in our experience, clients (HR teams) will ask technology vendors around AI – but they don't really help get to the crucial stuff...

1. Do you apply Deep Learning or Machine Learning?

Some people consider this to be an interesting question, but neither Machine Learning or Deep Learning is inherently better – as with most things, it depends on the application. The question to ask is: does the AI solve a relevant problem in an unbiased way? Whatever method is used, can we explain how a recommendation was derived?

2. What models are used?

Will it make a difference? It would be better (as we suggest above) to ask why a particular model has been chosen in each case. And, ultimately, why AI is being used, rather than other methods. What value does it drive?

3. What is the size of the dataset?

Sure, this could imply a more balanced data set, and more balanced training models. But that is the important question! How representative of the real world is the data set?

4. What is the size of your data science team?

Are you asking your vendors about how many people are working on AI in their organization? Alas, this doesn't tell you if the models are better (and team quality is more important than size). Perhaps ask: What are they working on? And what outcomes are you trying to drive?

5. Do you have a chatbot?

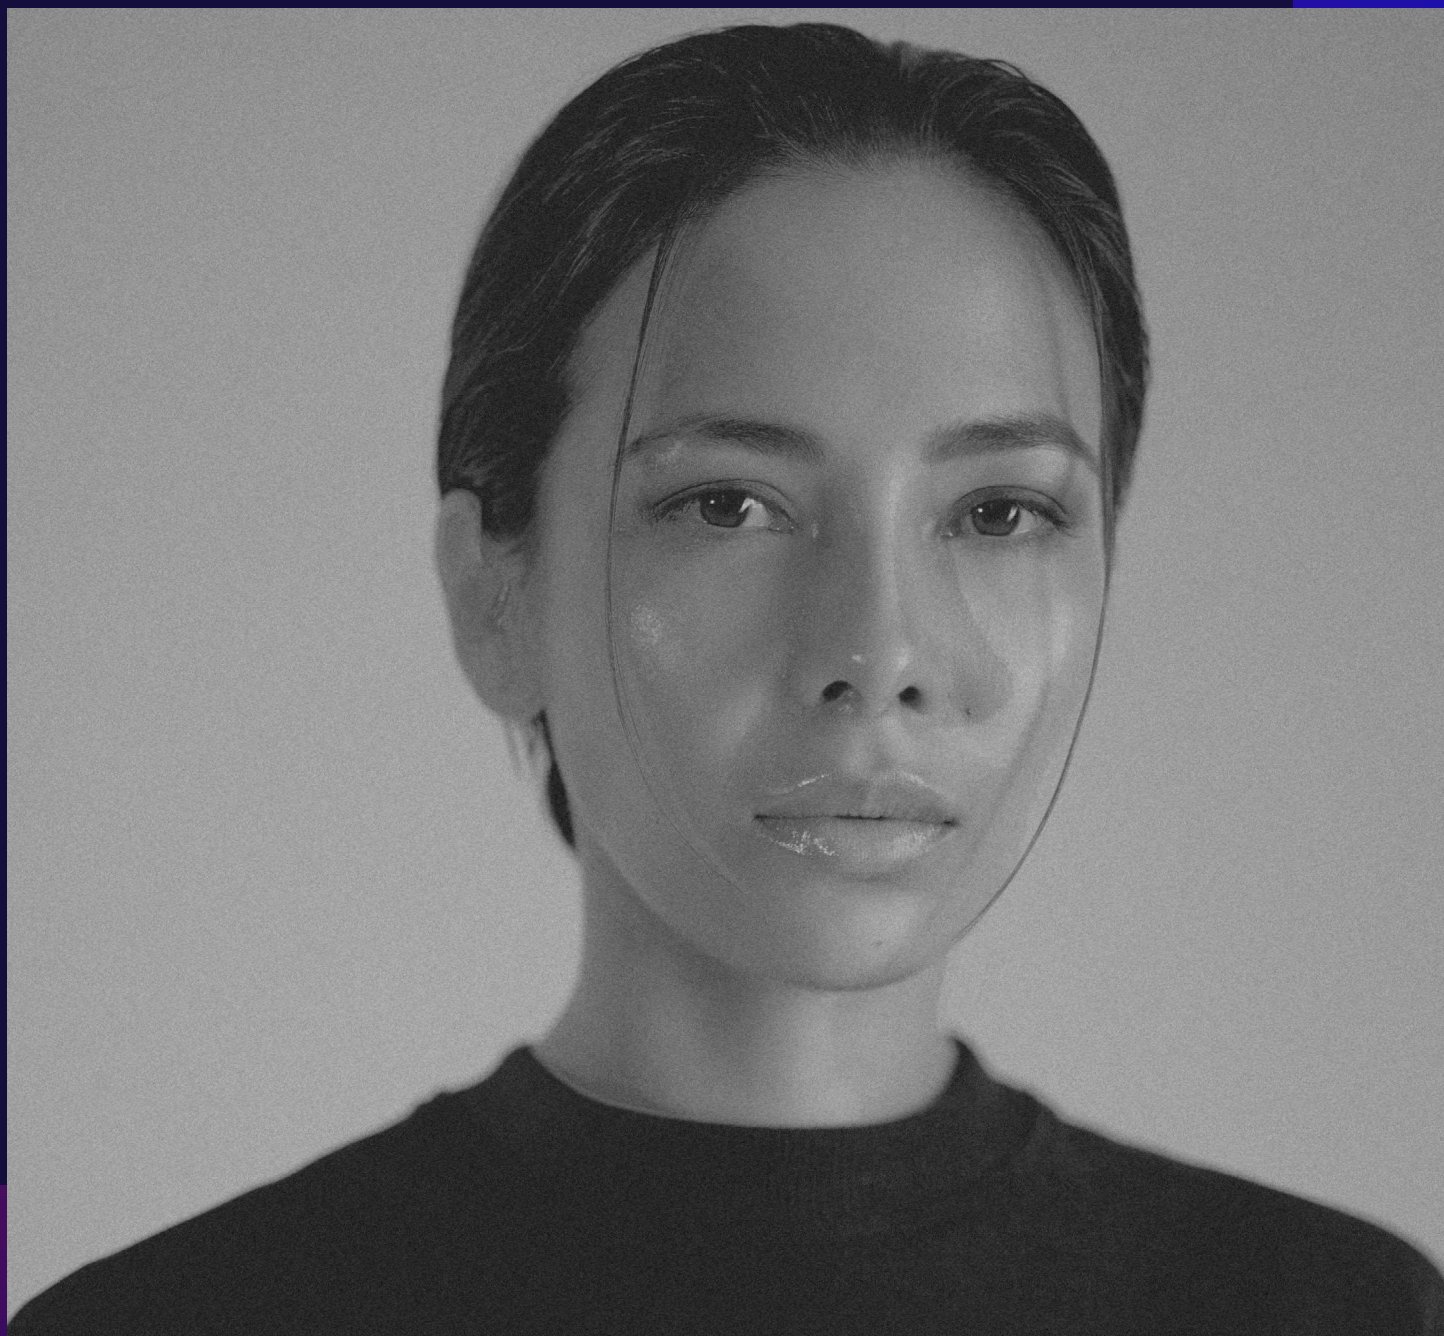
For some reason, this seems to be a common question from HR teams looking to explore AI. Of course, a chatbot is not always the answer. And not all chatbots are powered by AI! The bigger question is: What is an ideal user/candidate experience, and do we need to use AI (and how) to deliver the highest quality experience?

Conclusion

Today, with growing demands on Talent teams and new regulations around AI in HR, it is more important than ever to understand the true value of any technology you implement, across the talent lifecycle.

Remember: AI is not a magic bullet. Check it is meeting your needs, and delivering value for HR teams. Look for quality data, explainable models, and valuable experiences. Is it improving things for users, candidates and employees, or just adding to the cognitive load? And can you be sure that it is helping make fairer decisions, and not adding bias to the process?

Learn more about how AI can be used in HR here, or get in touch with Beamery today for information on our own explainable AI and AI-powered solutions.





This document and all its contents is the property of Beamery Inc. THIS DOCUMENT IS FOR INFORMATION PURPOSES ONLY AND DOES NOT CONSTITUTE LEGAL ADVICE. DIAGRAMS, PERCENTAGES AND/OR OUTCOMES USED IN THIS DOCUMENT ARE FOR ILLUSTRATION PURPOSES ONLY. INDIVIDUAL OUTCOMES VARY BY CUSTOMER. THE BEAMERY CUSTOMER IS SOLELY RESPONSIBLE FOR ITS USE OF THE BEAMERY TECHNOLOGY. BEAMERY WILL NOT BE RESPONSIBLE FOR ANY LIABILITY RESULTING FROM OR CONNECTED WITH THE USE OF THIS DOCUMENT.

Beamery Ltd 2023. All rights reserved.