

March 2025

## How HireVue Created “Glass Box” Transparency for its AI Application

Monideepa Tarafdar

Irina Rets

Lindsey Zuloaga

Nathan Mondragon

Follow this and additional works at: <https://aisel.aisnet.org/misqe>

---

### Recommended Citation

Tarafdar, Monideepa; Rets, Irina; Zuloaga, Lindsey; and Mondragon, Nathan (2025) "How HireVue Created “Glass Box” Transparency for its AI Application," *MIS Quarterly Executive*: Vol. 24: Iss. 1, Article 7.  
Available at: <https://aisel.aisnet.org/misqe/vol24/iss1/7>

This material is brought to you by the AIS Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in MIS Quarterly Executive by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact [elibrary@aisnet.org](mailto:elibrary@aisnet.org).

## How HireVue Created “Glass Box” Transparency for its AI Application

*HireVue develops and sells an AI application for assessing job interviews. We describe several transparency-related challenges HireVue faced and explain how it addressed them through five types of transparency practices. Drawing from these practices, we provide recommendations for AI software companies to create an AI transparency. Instead of focusing on opening a technical “black box” to disclose how an AI algorithm works, transparency should encompass software design and development processes, as well as organizational functions, AI application interfaces, clients’ needs and regulatory requirements.<sup>1,2</sup>*

**Monideepa Tarafdar**  
University of Massachusetts Amherst (U.S.)

**Irina Rets**  
The Open University (U.K.)

**Lindsey Zuloaga**  
HireVue (U.S.)

**Nathan Mondragon**  
HireVue (U.S.)

### Learning from HireVue’s Approach to AI Transparency

In this article, we describe HireVue’s “glass box” approach to AI transparency. (Our research method is described in the Appendix.) HireVue develops, sells and deploys an AI hiring application that analyzes video interviews of job applicants to predict an interviewee’s (future) success in that job. It operates across industries (including finance, retail, healthcare and the public sector) and works with over 1,160 clients in more than 100 countries. At the time of writing, the organization had over 350 employees. Established in 2004 in Salt Lake City, Utah, it was the earliest entrant in the global AI hiring market.

HireVue’s clients source the applicants who complete video job interviews by answering questions designed to assess competencies needed for the listed job. The competencies and the interview questions to measure them are drawn from a library developed by HireVue, based on its research into different jobs, consultations with its clients and publicly available standardized occupation-specific and job-specific competencies from the Occupational Information Network



<sup>1</sup> Fred Niederman is the senior accepting editor for this article.

<sup>2</sup> We thank Professor Niederman for his invaluable guidance in helping us shape and improve this article and thank the anonymous reviewers for their constructive feedback. Our research was funded by grants from the U.K.’s Economic and Social Research Council (ESRC), award number ES/T012382/1, and the Social Sciences and Humanities Research Council of Canada (SSHRC) as part of its Canada-UK Artificial Intelligence Initiative, award number SSHRC 2003-2019-0003.

(O\*NET) database.<sup>3</sup> AI algorithms are used to predict the competencies from the answers to the questions.

The algorithms employ supervised machine learning, based on past interview responses, to predict an applicant's rating on the competency in question. Over the past 20 years, HireVue has conducted more than 66 million interviews of applicants for clients who use its AI hiring application. A subset of the interviews is used to train the AI algorithms for competency measurements.<sup>4</sup>

HireVue's experts in occupational psychology rate these interviews for the competencies. The raters are a diverse group (by race, sex, age) and undergo training on rating job competencies to ensure their ratings are consistent and accurate. Interrater calibration reviews are conducted and ratings are evaluated for race, sex and age differences before they are used in the algorithms as training data. The AI algorithms are tested and validated to measure the competencies at accepted levels of psychometric validity (e.g.,  $r > 0.60$ ).<sup>5</sup> They are also tested for and mitigated against statistically significant group differences

or adverse impact (e.g., differences in scores attributed to sex, age, race).<sup>6</sup>

Hiring-related AI policy and legislation focuses primarily on privacy and potential harms such as bias. The EU AI Act classifies AI hiring applications as high risk because they have a high potential for doing harm.<sup>7</sup> These applications impact not only a person's right to work but also the right to equality and non-discrimination. They have a long reach because they affect people's life prospects, including what work they will do, how much they will get paid, where they will live and the quality of life they will have. In turn, hiring applications significantly impact a person's sense of autonomy, self-identity and active membership in society. In short, they can determine who gets access to the economic opportunities the world offers.<sup>8</sup>

Thus, the EU AI Act includes requirements for data quality and non-discrimination. AI software companies in the hiring sector must provide technical documentation as evidence of compliance, risk management and human oversight over AI use. Moreover, the EU's General Data Protection Regulation (GDPR) requires AI software companies to describe their data protection policies.<sup>9</sup> In the U.S., the Algorithmic Accountability Act requires companies to assess their AI-driven systems “for impacts on accuracy, fairness, bias, discrimination, privacy, and security.”

HireVue innovated approaches and techniques for developing AI-enabled software, taking account of these regulatory requirements for the highly sensitive nature of hiring applications. Though transparency is important for all AI

3 For information about O\*NET, see Dye, D. and Silver, M. “The Origins of O\*NET,” in Peterson, N. G., Mumford, M. D., Borman, W. C., Jeanneret, P. R. and Fleishman E. A. (eds.) *An Occupational Information System for the 21st Century: The Development of O\*NET*, American Psychological Association, 1999, pp. 9-19.

4 The video interviews are conducted concurrently across many applicants. For example, 100 applicants can complete video interviews at the same time every hour of every day. High volumes of interviews are conducted outside of working hours (evenings and weekends). HireVue pulls samples of interviews with the right competency questions to use for training in the AI development process. Consent to use the interview is part of the terms and conditions the applicant agrees to when initiating the interview. Thus, a much smaller number of interviews (around 45,000) are pulled for training. Additionally, the reliability and validity of the AI models are maximized with the sampled interviews used for training. In psychometric testing, the  $r$  value is a measure of the relationship between a test and a criterion, such as job performance. It is reported as a validity coefficient, which is a number between 0 and 1. For a detailed review of how HireVue's AI algorithms were developed and tested, see Liff, J., Mondragon, N., Gardner, C., Hartwell, C. J. and Bradshaw, A. “Psychometric Properties of Automated Video Interview Competency Assessments,” *Journal of Applied Psychology* (109:6), June 2024, pp. 921-948, available at <https://doi.org/10.1037/apl0001173>.

5 In psychometric testing, the  $r$  value is a measure of the relationship between a test and a criterion, such as job performance. It is reported as a validity coefficient, which is a number between 0 and 1. For a detailed review of how HireVue's AI algorithms were developed and tested, see Liff, J., Mondragon, N., Gardner, C., Hartwell, C. J. and Bradshaw, A. “Psychometric Properties of Automated Video Interview Competency Assessments,” *Journal of Applied Psychology* (109:6), June 2024, pp. 921-948, available at <https://doi.org/10.1037/apl0001173>.

6 Rottman, C., Gardner, C., Liff, J., Mondragon, N. and Zuloaga, L. “New Strategies for Addressing the Diversity-Validity Dilemma with Big Data,” *Journal of Applied Psychology*, (108:9), September 2023, pp. 1425-1444, available at <https://doi.org/10.1037/apl0001084>.

7 See *EU AI Act: First Regulation on Artificial Intelligence*, European Parliament, June 8, 2023, available at <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>. High-risk applications include those used in hiring, credit checking and the justice system. The proposed rules would require such AI applications to use high-quality data, document data provenance and traceability, and account for human oversight.

8 Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K. “Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices,” in *Proceedings of FAT\* '20: Conference on Fairness, Accountability, and Transparency*, Barcelona, Spain, January 2020, pp. 469-481.

9 Chae, Y. “U.S. AI Regulation Guide: Legislative Overview and Practical Considerations,” *The Journal of Robotics, Artificial Intelligence & Law* (3:1), 2020, pp. 22.

applications, it is especially so for areas like hiring. To address the transparency challenges it faced, HireVue created an AI transparency “glass box” rather than opening up a technical “black box.” The HireVue case therefore serves as an example of the approach that AI software companies can use to provide transparency for clients who use AI applications.

## HireVue’s AI Development and Deployment Process

HireVue’s AI hiring application assesses the text of each applicant’s interview for the job competencies, using natural language processing (NLP) techniques. The outcome of the AI-based assessment is a score assigned to the applicant that indicates how well the applicant will likely perform in the key competencies required for the advertised job. The client then makes the final selection of applicants. The AI hiring application can be integrated with clients’ applicant tracking systems.

HireVue uses the following five-step process for developing and deploying the AI hiring application for the majority of its clients:

1. *Job analysis:* In the first step, the competencies needed for a given job are identified primarily by HireVue’s occupational psychology department in consultation with the client. Examples of job-related competencies include “Willingness to Learn,” “Adaptability” and “Customer Orientation.”
2. *Assessment design and validation:* In Step 2, the identified competencies are matched with HireVue’s library of competencies and questions, employing the previously developed AI algorithms to evaluate the competencies through answers to the questions.
3. *Assessment deployment:* Once the final set of interview questions has been identified and agreed upon by HireVue and the client, HireVue then loads the questions to its interview platform. Applicants selected by the client complete their job interviews by verbally answering the questions through the platform’s online video interviewing software. HireVue transcribes the answers and uses them as input for its AI

hiring application to predict applicants’ competencies based on their interview responses.

4. *Post-launch assessment validation:* This step includes analyses of bias and predictive accuracy, performed after the AI hiring application has been used by the client for some time, typically 9 months to a year. The goal is to examine if the application is providing client value by identifying better applicants, as measured by retention and other job performance metrics.
5. *Feedback:* In the final step, the results from Step 4 are communicated to the client for analysis of future jobs.

## Evaluating the Benefits of the AI Hiring Application

HireVue evaluates the client benefits of its AI hiring application along two dimensions. The first measures the time saved by a client’s hiring manager and improvements in fairness outcomes of the client’s hiring process. The second measures improvements in the quality of employees who were hired after being interviewed and assessed by the AI hiring application. These measures are based on reduced turnover and better job performance (e.g., higher sales, improved call center metrics, higher customer satisfaction).

Additionally, HireVue also evaluates the application against its own business metrics, such as improved client retention and development of and compliance with its own AI fairness methods. The latter includes reducing adverse impact outcomes compared to traditional or non-AI assessment methods and publications in leading peer-reviewed occupational psychology journals.

## The AI Transparency Challenges Faced by HireVue

HireVue encountered four main challenges in providing transparency for its AI hiring application:

1. *Transparency means different things to different roles.* For example, transparency for a data scientist includes disclosure of algorithmic training and validation details or mathematical techniques for

extracting model features and their respective weights. However, that level of detail would be meaningless for most HireVue clients, the exception being the few technically minded ones. Most clients prefer just enough information to understand the AI models’ predictions and accuracy. Job applicants want information on what is being measured and assurance that they won’t be discriminated against by the AI application. With these stakeholders having such different expectations for transparency, giving the wrong kind of information would be risky. For example, giving detailed technical information to clients who did not want it could create unnecessary confusion and could even damage trust if misinterpreted.

2. The second challenge is the difficulty of disclosing all these different types of information because of *sensitivities about sharing the AI application’s intellectual property*. For example, explaining to job applicants what is being measured by the AI hiring application or giving out key technical details to clients might make HireVue vulnerable to competitors. Many AI software companies refrain from sharing such details of their products with external parties. Others carefully tread a line between sharing and making themselves more vulnerable to competition.
3. Without a reasonable degree of technical understanding of the application, clients may find it challenging to *properly interpret the AI application’s outputs* and how the outputs should be used in conjunction with human agency. Clients also may find it challenging to *understand how the AI application needs to be monitored, maintained and changed*. Even though the AI algorithms may be “locked” after deployment, continuous monitoring is required from a client because the algorithms may behave differently with different populations and evolve as the data and usage context changes.
4. The fourth challenge arises because the AI hiring application, once deployed, often comes up with *unexpected and undesirable*

*outcomes that could not be predicted beforehand*, such as biased scoring of applicants’ competencies based on their personal attributes. Clients can become confused and alarmed by such outcomes.

## Five Types of Practices for Addressing the AI Transparency Challenges

We categorized the practices HireVue used to address the AI transparency challenges into five types—pre-deployment client-focused practices, internal practices, post-deployment client-focused practices, knowledge-related practices, and audit- and industry-related practices. Together, these five types of practices encompass activities across all organizational departments. Given the essentially probabilistic and technically black-boxed nature of AI algorithms, there are limits to the degree to which disclosing an algorithm’s parameters, weights, computations and data is possible because the outputs are inherently dynamic and unpredictable. In essence, HireVue’s AI hiring application is a “black box” because of the nature of the algorithm’s computational models (e.g., large language models and neural networks), making its inner workings difficult to understand. To compensate for this, HireVue’s approach was to open up the *sociotechnical* components associated with the application, such as data provenance, training strategies, and bias-related analysis and interpretation of outputs. This approach enabled it to reveal to clients, applicants, auditors, regulatory bodies and its own employees, the different activities associated with the application’s design, deployment and use.<sup>10</sup> The five practice types are depicted in Figure 1 and described in detail below.

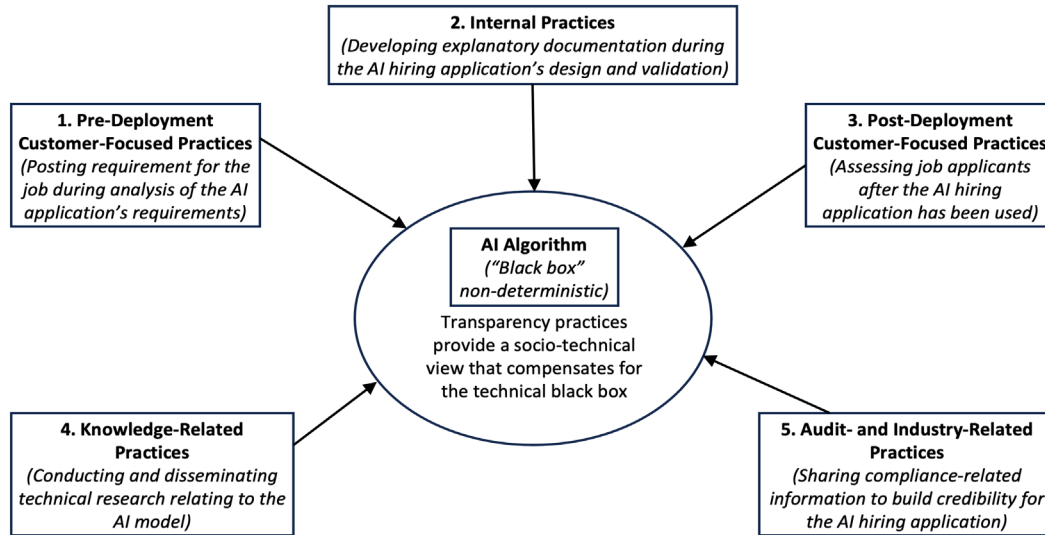
### 1: Pre-Deployment Transparency-Focused Client Interaction Practices

Before deploying the AI hiring application on a client, HireVue is open about the applicability and utility of AI-based hiring assessments. Product

<sup>10</sup> This approach can be likened to the concept of an open system that has a technical and relatively black-boxed core but that interacts with developers, users, clients and auditors through open sociotechnical components. The purpose of the transparency practices is to create open sociotechnical components. For more information, see von Bertalanffy, L. *General System Theory*, George Braziller Inc., 1969.



Figure 1: HireVue’s Five Types of Transparency Practices



managers and solution architects have frank conversations with clients as to whether AI hiring is needed for a given job. They explain that the AI-based assessment will work only when the volume of applicants is large. As HireVue’s CEO explained:

*“If, for example, the client says, ‘I’m interviewing so many candidates. I recognize that as an evaluator I can’t watch 1,000 interviews with any consistency. And so how do I find tools that help with that,’ [we tell the client] AI works [in this situation. However,] if they are evaluating five people, I wouldn’t recommend AI assessments.”*

Pre-deployment, product managers also explain to clients the boundary of influence of the AI application. They specify that the application’s output is limited to a categorization of client-selected applicants based on predictions of the extent to which they have the competencies needed for the job. The eventual hiring decisions will be made by the client and are influenced by where the client advertises the job, how it selects applicants for the interviews, and who they eventually choose based on the application’s output. As explained by a HireVue product manager, these conversations help the client understand

*“How they can use the AI output, what it can do for them and what it can’t. We never say make 100% of your selection decisions based on an algorithm. We help them understand that the AI assessment is part of their strategic hiring process, and not 100% of it.”*

HireVue encourages clients to clearly explain to applicants that their interviews would be evaluated by an AI application and to include messaging that encourages applicants to be authentic and tell their own stories about their experiences in their answers.

Clients tend to have many pre-deployment questions about the AI application. For example, they may want to know if an applicant who spoke quickly without verbal pauses or did not look into the camera would score low. To address this concern HireVue’s occupational psychologists clarify what parameters and variables in the AI hiring application are relevant for prediction. As recounted by an occupational psychologist, they explain to clients that the AI application “doesn’t care about any of those things. What we care about is the competencies that they are exhibiting with their responses.”<sup>11</sup>

<sup>11</sup> HireVue tested and validated that speech and facial movement variables do not add incremental validity when predicting job success beyond text analyses.

**Table 1: HireVue’s Practices for Addressing Pre-Deployment Client-Focused Transparency Challenges**

Transparency Practice	How the Practice Contributes to Transparency	Transparency Challenge Addressed by the Practice
Product managers explain to clients the applicability and utility of AI-based hiring assessments.	<ul style="list-style-type: none"> <li>• Helps clients understand the AI hiring application and when it is useful.</li> <li>• Clarifies clients’ misconceptions about AI-based hiring such as AI being the panacea for all hiring problems or that AI hiring is always biased and unfair.</li> </ul>	Transparency means different things to different roles—most clients want just enough information to understand what the AI application would predict and how accurately.
Occupational psychologists explain to clients what parameters and variables are included in the AI hiring application.	<ul style="list-style-type: none"> <li>• Helps clients understand the AI hiring application.</li> </ul>	Clients find it difficult to properly interpret the AI hiring application’s outputs without a reasonable degree of technical understanding of the application.
Product managers explain the potential pitfalls of the AI hiring application, such as bias.	<ul style="list-style-type: none"> <li>• Helps clients understand the complexity of the AI hiring application.</li> <li>• Manages clients’ expectations.</li> </ul>	The AI hiring application, once deployed, often reveals unexpected and undesirable outcomes that could not be predicted beforehand.

HireVue also developed and published an “AI Explainability Statement”<sup>12</sup> for clients, that describes in general language the computations and factors underpinning the AI algorithms. Given both the positive and negative aspects of AI, product managers also describe to clients the potential pitfalls of AI models, such as bias.

Table 1 summarizes how the Type 1 practices address the pre-deployment transparency challenges.

## 2: Internal Transparency Practices

In anticipation of clients’ questions, HireVue’s data scientists documented the processes behind the AI application’s design and training. A senior data scientist said that clients had “all sorts of questions about the AI models. A lot of our work is trying to make sure that we have thought about these questions and that we have good answers for these questions.” The AI models are complex and include over 5,000 linguistic features that are analyzed and matched to competencies in a single interview transcript. However, most clients are not interested in the inner workings of the AI models but want to know how the AI application would make their hiring decisions

faster and more efficient and improve applicants’ experiences. Thus, it is important to provide explanations about topics that are important to clients rather than sharing technical papers likely to overwhelm them with model details.

HireVue therefore developed extensive but easy-to-understand explanations for clients and delivers them on an as-needed basis. A HireVue solution architect explained that if clients do want to know about the AI model and training,

*“Then, of course, we will go into all the details. If clients are interested, we have 50-60-page documents that we will share with them ... validation papers and all of the research that we have done. But, ... it may not be something we would offer upfront. You don’t want to necessarily confuse the audience by saying, ‘Hey, nice to meet you, here’s a 50-page document about AI; read that and then we’ll talk afterwards.’”*

For clients who are interested, technical explanations cover AI model design and training, helping them to understand the importance of both domain/subject (occupational psychology) and technical (machine learning/statistics) knowledge in the AI application and appreciate its complex nature.

<sup>12</sup> The statement is available at <https://www.hirevue.com/legal/ai-explainability-statement>.

**Table 2: HireVue’s Practices for Addressing Internal Transparency Challenges**

Transparency Practice	How the Practice Contributes to Transparency	Transparency Challenge Addressed by the Practice
Data scientists document the processes behind the AI application’s design and training.	<ul style="list-style-type: none"> <li>Formulates answers to potential questions from clients about how the AI application was developed.</li> </ul>	Clients find it difficult to properly interpret the AI application’s outputs without a reasonable degree of technical understanding of the application.
Solution architects develop extensive but easy-to-understand explanations for clients and deliver them on an as-needed basis.	<ul style="list-style-type: none"> <li>Helps clients understand the complexity of the AI application.</li> <li>Addresses information needs for different clients based on preferences regarding how much detail of information they want. Many clients do not want to be overwhelmed with a lot of technical information.</li> </ul>	Transparency means different things to different roles—clients want just enough information to understand what the AI model would predict and how accurately. Clients find it difficult to properly interpret the AI application’s outputs without a reasonable degree of technical understanding of the application.
Multiple departments contribute to the explanations.	<ul style="list-style-type: none"> <li>The client has access to a wide range of information spanning multiple departments’ views of different aspects of the AI hiring application.</li> </ul>	Transparency means different things to different roles.

Multiple HireVue departments contributed to these explanations, depending on the nature of the information. While occupational psychologists provided definitions and examples of competency behaviors, data scientists explained, from an NLP viewpoint, how specific words and phrases relate to particular competencies. An occupational psychology consultant explained that to support a client who had many technical questions about the models, she was able to “tap on the shoulder” of a data science colleague specializing in assessment psychometrics.

Table 2 summarizes how the Type 2 practices address the internal transparency challenges.

### 3: Post-Deployment Client-Focused Transparency Practices

The output of the AI hiring application is the automatic scoring of applicants’ interviews for each competency identified as necessary for the job. The application displays all applicants’ scores in a ranked list split into three bands: bottom, middle and top third, based on predictions of their job success. Figure 2 shows an example of the three-tier banding of applicants produced by the AI hiring application.

The three-tier ranking prevents the client’s hiring manager responsible for assessing

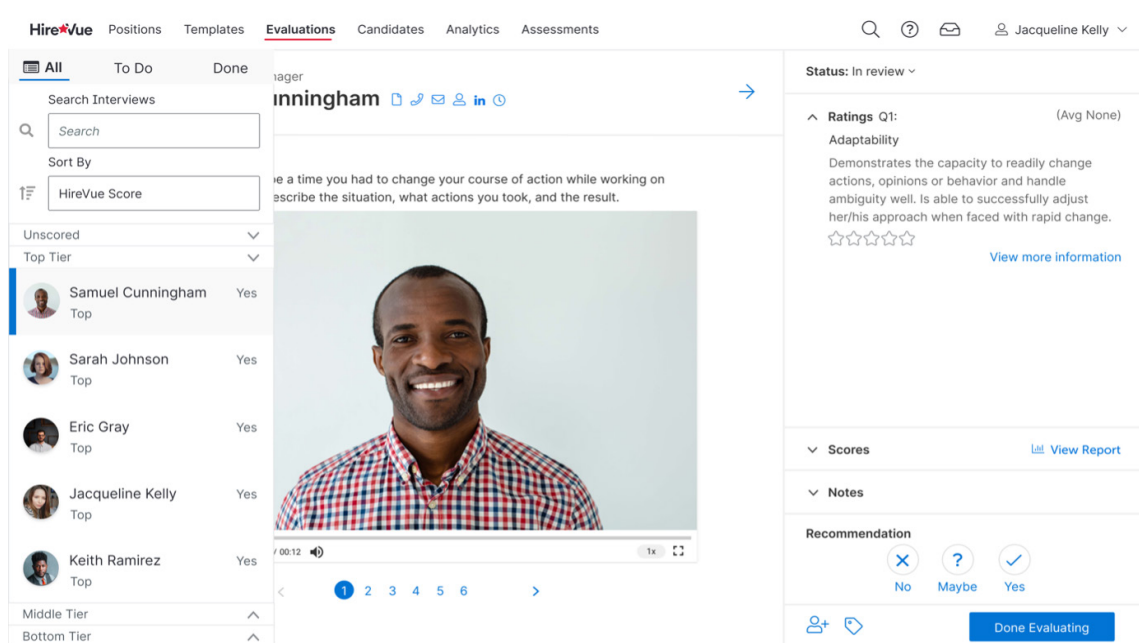
applicants from being overwhelmed by small differences between applicants (e.g., score of 85 vs. 89 percentiles) and helps the manager to evaluate all applicants. HireVue’s CEO explained that “the difference between a 95 and a 92 was not that big,” and that HireVue tries “to get clients away from believing that [the AI assessment was that finely graded] that you could tell a 90 from a 91.” As the company’s chief occupational psychologist explained: “So, we don’t want the client to [mistakenly] think ‘Well, I want the 95, not the 92, because it is three levels bigger.’ We will band the candidates into, bottom, middle and top third [based on percentile values].” The three-tier banding makes it easy for clients’ hiring managers to interpret the AI application’s outputs in a practical, meaningful and easy-to-understand way, as opposed to a raw percentile ranking.

The AI application’s output also has the option for a more in-depth report, in the form of explanations accompanying the AI scores that show what competencies are being measured and how each applicant scored against each of them and a video link to the interview. Figure 3 shows a sample of the drill-down output providing details of an applicant’s interview performance.

The combination of these AI output options enables clients to use the outputs flexibly,



Figure 2: Three-Tier Banding of Applicants’ Interview Performance



as explained by HireVue’s chief occupational psychologist: “Sometimes they will use it as a guide, and they will still watch all the videos, but they will watch the red ones with less intensity, for example.” As a consequence, HireVue can explain to clients that it is the client and not the AI application’s score that makes the final hiring decision and that clients can choose how to use the AI application’s score for the applicant.

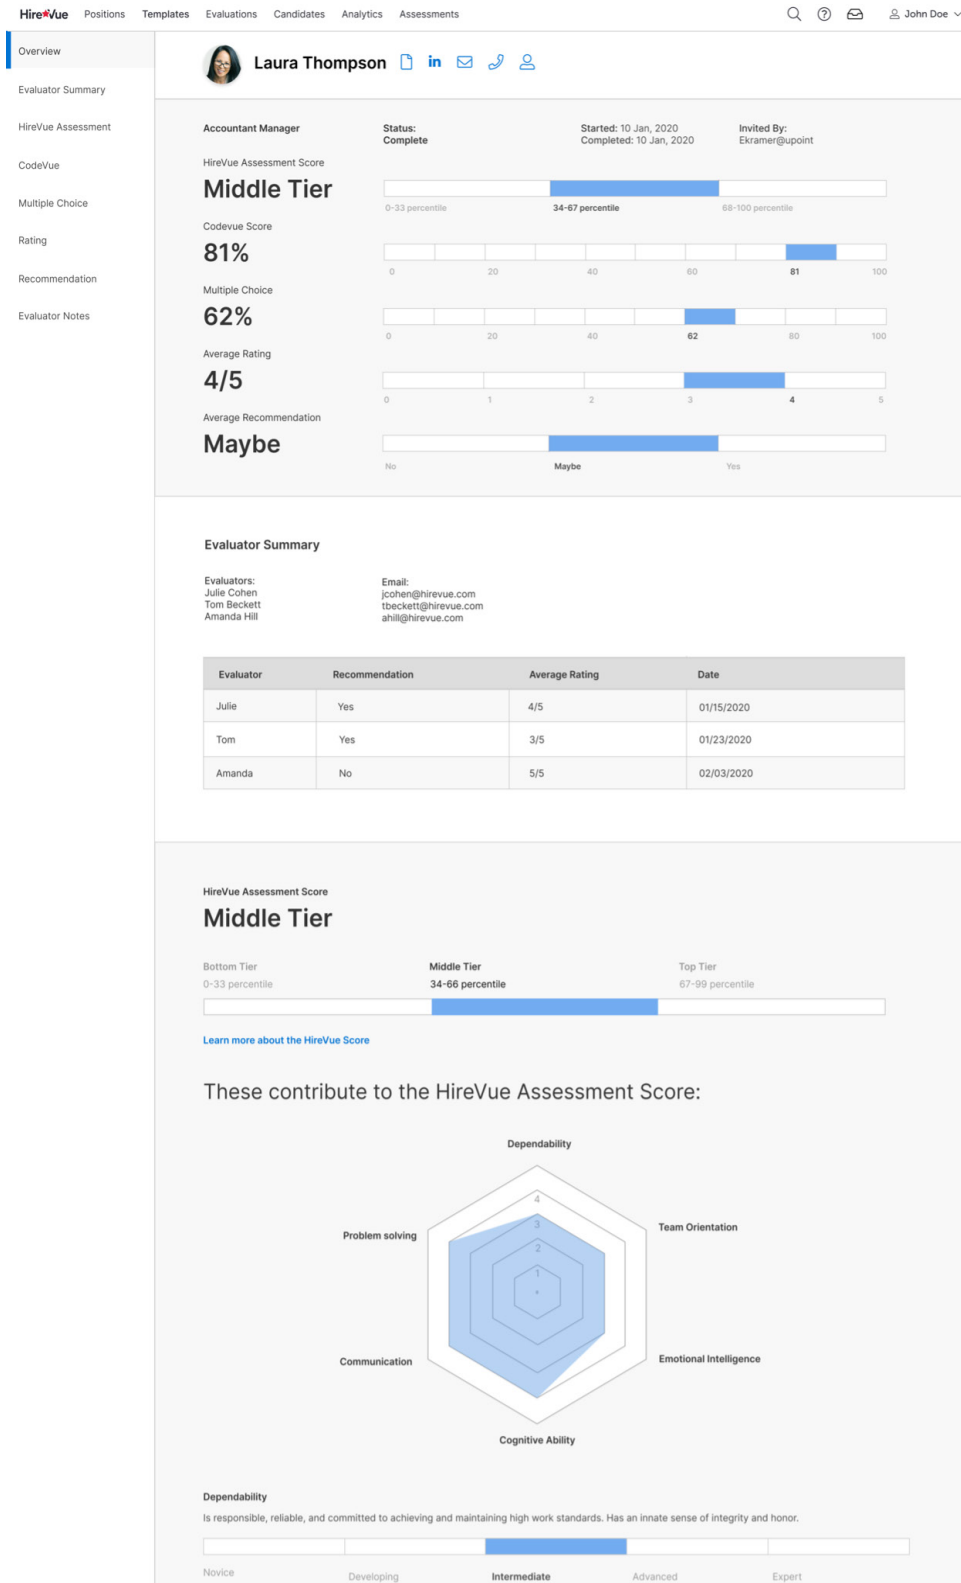
Another post-deployment transparency practice used by HireVue is that it engages with clients when they disagree with the AI application’s outputs. A HireVue occupational psychology consultant stated:

*“Clients have come to us and said, ‘My recruiters are scoring candidates very differently than your model, we’re not trusting the model.’ We respond by saying, ‘let’s see how you’re scoring communication on a candidate, let’s walk through a couple of samples, let’s see how the model’s scoring it, let’s look at the behaviors that need to be elicited to land on a certain proficiency score and let’s see what’s different.’ So, it*

*calibrates and it gets everybody on the same page. And I can tell you when I’ve done that they turn around and go, ‘Oh, OK, I get this now.’”*

HireVue also collaborates with clients to analyze and validate the AI application’s outputs after deployment, typically between 9 months and a year later. Validation involves reviewing the scores of selected applicants for bias and prediction power, as described by the company’s chief occupational psychologist: “[We want to see] if, now that they are employees, the people who scored higher [in their job application interviews] are performing higher and if the people who scored lower are performing lower.” For example, if applicants hired from minority backgrounds are rated by the client as poor performers on the job, the rating could be due to workplace culture and the resources available to the employee. In the words of HireVue’s chief occupational psychologist, “Sometimes it is the case that we clean the problem [bias] at the assessment stage, and then the client dirties it up once the candidates become employees.” HireVue

Figure 3: Sample Drill-Down Features of an Applicant’s Interview



**Table 3: HireVue’s Practices for Addressing Post-Deployment Client-Focused Transparency Challenges**

Transparency Practice	How the Practice Contributes to Transparency	Transparency Challenge Addressed by the Practice
Display all applicants’ scores, using a three-tier relative bottom, middle and top-third banding calculation based on prediction of their job success.	<ul style="list-style-type: none"> <li>Provides clients with a practical, meaningful and easy-to-understand way to look at the AI hiring application’s outputs.</li> </ul>	Clients find it difficult to properly interpret the AI application’s outputs without a reasonable degree of technical understanding of the application.
Provide in-depth explanations to accompany the AI scores.	<ul style="list-style-type: none"> <li>Explains to clients the logic behind the AI application’s outputs, equipping them better to make selection decisions.</li> <li>Enables clients to examine the outputs in as much detail as they need to.</li> </ul>	Clients find it difficult to properly interpret the AI application’s outputs without a reasonable degree of technical understanding of the application. Transparency means different things to different roles—clients want information on an as-needed basis.
Engage with clients when they disagree with the AI application’s outputs.	<ul style="list-style-type: none"> <li>Helps clients understand the complexity of the AI hiring application by explaining why and how the outputs come about.</li> </ul>	Clients find it difficult to properly interpret the AI application’s outputs without a reasonable degree of technical understanding of the application.
Collaborate with clients to analyze and validate the AI application’s outputs after deployment.	<ul style="list-style-type: none"> <li>Helps clients understand where things are not working by identifying adverse impacts and biases.</li> </ul>	The AI hiring application, once deployed, often reveals unexpected and undesirable outcomes that could not be predicted beforehand.
Communicate with clients to explain AI model changes.	<ul style="list-style-type: none"> <li>Helps clients understand how the AI application has been improved, based on research on emerging technology advancements.</li> </ul>	Clients find it difficult to understand how the AI application needs to be monitored, maintained and changed without a reasonable degree of technical understanding of the application.

communicates such learnings to the client, alerting them to problems and working with them to unearth potential causes. The organization’s goal is to continue collaboration with clients for the annual validation of the AI hiring application.

Post-deployment, HireVue also communicates with clients to explain the rationale for AI model changes. One such change was a switch from developing competencies for each client’s job description to selecting competencies from a pre-developed and existing library for a given job description. The first method involved capturing 300 or more completed interviews from the client’s existing employees and benchmarking against their performance in their roles to identify competencies to be evaluated for future applicants applying for similar job roles. The second method uses a large number

of interviews from HireVue’s own database of scored interviews that represents a variety of demographic and geodiverse populations, with (according to a HireVue principal solutions architect), “several thousand candidate responses against similar interview questions.” To build this library, HireVue used multiple trained expert raters to assess how past applicants answered the same interview questions on a behaviorally anchored rating scale.

This second assessment approach is more modular and means that AI model training is not limited to only using client data; instead, there is a much larger and demographically representative number of interviews to choose from. HireVue explained to clients that the earlier approach ran a greater risk of introducing biases from the client’s past performance data into the

**Table 4: HireVue’s Practices for Addressing Knowledge-Related Transparency Challenges**

Transparency Practice	How the Practice Contributes to Transparency	Transparency Challenge Addressed by the Practice
HireVue’s psychologists and data scientists conduct research and enhance their knowledge of AI interviewing and assessment models.	<ul style="list-style-type: none"> <li>Keeps their expertise updated so that state-of-art knowledge is applied to the AI application, spanning as much emerging knowledge as possible.</li> <li>Data scientists’ view of computational/mathematical techniques is incorporated into the AI models.</li> </ul>	Transparency means different things to different roles.
The psychologists and data scientists publish their research in peer-reviewed journals and as white papers	<ul style="list-style-type: none"> <li>Provides potential clients with detailed technical explanations of the AI hiring application without compromising intellectual property and competitive advantage.</li> </ul>	Sensitivities around sharing the AI application’s intellectual property.

predictions, as described by HireVue’s CEO: “The client’s existing process will replicate the biases and problems that are in their existing process.”

Table 3 summarizes how the Type 3 practices address the post-deployment client-focused transparency challenges.

#### 4: Knowledge-Related Transparency Practices

Keeping up with new technologies and industry and academic research guides HireVue’s ongoing development of its AI hiring application. For example, in addition to text analysis, it initially used video facial features (to infer emotions from the applicants’ frowning or smiling) and audio features (such as intonation) to predict competencies. This led to negative press publicity. Research on emerging NLP techniques by the data science team led HireVue to conclude that facial and audio features did not add substantially to the predictive accuracy of the AI model. Moreover, including these features could potentially exacerbate bias. As HireVue’s chief occupational psychologist explained: “Research that was coming out in the scientific literature was that emotional states [differ] across cultures, and the tone of skin is hard to measure, leading to more error.” Subsequent versions of the application dropped the video features and the output included only the text of the applicant’s interview.

HireVue’s psychologists and data scientists collaborate with academic researchers to conduct research and enhance their knowledge of AI

interviewing and assessment models. The chief occupational psychologist quoted the example of a “leading academic researcher in AI technology-driven assessment” who is a sabbatical fellow with HireVue.

The psychologists and data scientists also publish their research in the form of white papers and peer-reviewed academic papers on AI hiring topics such as online structured interviewing,<sup>13</sup> validating and testing assessment approaches and models, and minimizing bias. This research helps to develop rigor in HireVue’s technical solutions and inform other researchers in the field about the concepts and theories it applies in the AI models, while also retaining control over the organization’s technical intellectual property. The white papers are available on the organization’s website and are used to provide potential clients more detailed technical explanations of the AI application.

Table 4 summarizes how the Type 4 practices address the knowledge-related transparency challenges.

13 See, for example: 1) Mondragon, N. J. “Artificial intelligence in automated scoring of video interviews,” in Kantrowitz, T., Reynolds, D. and Scott, J. (eds.) *Talent Assessment: Embracing Innovation and Mitigating Risk in the Digital Age*, Oxford University Press, 2023; 2) Rottman, C., Gardner, C., Liff, J., Mondragon, N. and Zuloaga, L. “New Strategies for Addressing the Diversity-Validity Dilemma With Big Data,” *Journal of Applied Psychology* (108:9), April 2023, pp. 1425-1444; and 3) Griswold, K., Phillips, J., Kim, M., Mondragon, N., Liff, J. and Gully S. “Global Differences in Applicant Reactions to Virtual Interview Synchronicity,” *The International Journal of Human Resource Management* (33:15), July 2021, pp. 2991-3018.

**Table 5: HireVue’s Practices for Addressing Audit- and Industry-Related Transparency Challenges**

Transparency Practice	How the Practice Contributes to Transparency	Transparency Challenge Addressed by the Practice
Independent third-party auditors audit the AI algorithm’s design, training and testing processes with regard to bias, fairness and regulatory compliance. Internal audits are conducted to identify adverse and fair outcomes.	<ul style="list-style-type: none"> <li>• Provides a way to communicate/ signal the accuracy of the AI algorithm without revealing the actual computational details (i.e., the intellectual property).</li> <li>• Develops credibility with clients, regulators and industry bodies.</li> <li>• Identifies unforeseen outcomes such as bias.</li> </ul>	Sensitivities around sharing the AI hiring application’s intellectual property may arise. The AI hiring application, once deployed, often reveals unexpected and undesirable outcomes that could not have been predicted beforehand.
Expert advisory and client advisory boards tap into industry and regulatory concerns.	<ul style="list-style-type: none"> <li>• Provides input on information that the boards expect HireVue to provide.</li> </ul>	Transparency means different things to different roles.

## 5: Audit- and Industry-Related Transparency Practices

HireVue engages independent third-party auditors to audit the design, training and testing processes of its AI models for specific use cases with regard to fairness, bias and regulatory compliance. The organization’s occupational psychology and data science units are also audited. The auditors’ reports are published on HireVue’s website. One such audit conducted by ORCAA (<https://orcaarisk.com/>) concluded that “the assessments work as advertised with regard to fairness and bias issues; [we] did not find any operational risks with respect to clients using them.”

HireVue also regularly performs its own internal audits on adverse impacts. Typically, AI applications are held to a higher standard of validity and fairness due to the scale at which they can operate. The standard of fairness in hiring is often the measure of “adverse impact” toward certain demographic groups. Adverse impact is described as the negative consequences on a particular group of people, often stemming from employment practices that appear neutral but actually have a discriminatory effect with respect to a protected attribute (e.g., race, sex, age). HireVue measures adverse impacts by various ratios (e.g., the four-fifths rule) and statistical group differences (e.g., the two standard deviations test). A HireVue occupational

psychology consultant explained that the company “was not required to do [internal audits]. Nobody’s telling us to do it. It’s really just the right thing to do.” The organization also has two external boards—an expert advisory board and a client advisory board—that enable it to tap into industry concerns that inform its practices regarding AI development and deployment.

Table 5 summarizes how the Type 5 practices address the audit- and industry-related transparency challenges.

## Recommendations for How AI Software Companies Can Create a Transparent AI “Glass Box”

AI applications affect people’s day-to-day lives by substantively influencing decisions—for example, who gets financial assistance, who has access to the labor market, who gets admission to what university and who receives what kind of judgments from the justice system. Yet, it is difficult for organizations (including companies, governments and courts) that use these applications to trace the computational and data provenance of the outputs on which such decisions are based. Those affected (e.g., job seekers, citizens, defendants) find it even more difficult because they do not even have access to the AI application. With transparency



deemed to be critical for AI-related legislation to be effective,<sup>14</sup> AI software companies and vendors will need to provide more than algorithmic transparency to clients, legal and regulatory entities, and other stakeholders.

The traditional concept of AI transparency focuses on opening a technical “black box” by disclosing how the AI algorithm works. Data scientists make the inputs, outputs and computational/statistical logic explicit where possible. However, it is not possible to fully explain or predict how the outputs of AI algorithms come about. Furthermore, such explanations become outdated as AI algorithms continue to learn from new production data. Instead of the traditional concept, we propose a deeper approach that focuses on transparency as the everyday work of making apparent the broader context of AI development and deployment.

Our view of AI transparency thus includes internal interactions among multiple departments, external interactions with clients, partners and regulators, developing technical knowledge and an understanding of the (un)intended effects of AI applications. Instead of seeing AI transparency as opening a black box, we perceive it as creating a “glass box” through which a range of actions and implications associated with the development, implementation and use of an AI application are revealed. Below, we provide five recommendations that AI software companies can use to create such a glass box.

## 1. Engage with Clients at Multiple Touch Points

HireVue uses client-focused transparency practices at multiple steps in the AI hiring application’s design and deployment. The focus of client interactions during pre-deployment job analysis is to clarify and explain the AI hiring application’s benefits (e.g., huge efficiencies in going through a large number of applications) and risks (e.g., regulatory risks vis-à-vis sex- or race-related bias and discrimination), why and when it should be used to screen applications and hire applicants and, broadly, how it works. Post-deployment client interactions focus on explaining the AI hiring application’s outputs

and working with the client to analyze and understand things such as bias with respect to applicants’ sex, race and other protected attributes.

HireVue needs these intensive client engagements because any inherent bias in its AI hiring application could lead to legal action by clients and applicants and to their faith in social equity and justice being undermined. For AI applications in less critical sectors (e.g., for purchasing or advertising), client engagement could be less intensive. We recommend that client engagement should vary in terms of *more* or *less* engagement-related actions over a *broader* or *narrower* range of the application’s design and deployment activities.

## 2. Design Interfaces that Enable Users to Interrogate AI Outputs

HireVue presents the AI application’s outputs in terms of a banded (bottom, middle and top third) depiction of applicants’ competency scores (a common practice in hiring assessment). This depiction serves two purposes. First, the banding is *suggestive* in that it helps clients to get a sense of applicants’ relative performance. Second, it gives clients a sense of agency—i.e., they can decide how to interpret each band. Clients could decide, for example, to start by evaluating top-tier applicants and move down the list. Or they could decide not to consider the bottom band at all. Thus, each client can interpret the AI application’s output tailored to its hiring decision-making process.

Suggestive representations such as banding highlight important aspects of AI applications’ outputs that users might pay attention to, increasing the efficiency and effectiveness of clients’ interpretation and decision-making. However, in addition to these suggestive representations, HireVue’s clients can also access more detailed quantitative and qualitative information by drilling down on each applicant’s scores. This feature creates opportunities for users to interrogate the AI output in a way that does not stifle thoughtfulness and innovation in their decision-making. For decisions that have critical societal consequences, such as medical triage, hiring and judicial cases, it is important that users are able to interrogate the outputs of AI applications and dig deeper into those aspects

<sup>14</sup> Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K., op. cit., January 2022.

that they want explanations for.<sup>15</sup> However, suggestive representations of AI outputs may suffice and be more efficient, for less critical and more commercial areas such as advertising.

Regardless of the type of decision, greater interrogation capability is costlier to develop and more time-consuming to use, an interesting trade-off that product developers should consider. The types of analysis and visualization for presenting the outputs can vary depending on clients’ preferences. Though HireVue used three-tier banding, alternative statistical representations such as clusters could also be used.

### 3. Communicate About AI Regulation Compliance Activities

Laws regulating AI differ by country, region and industry sector. The EU AI Act classifies AI applications into three risk categories. AI hiring applications that rank applicants are in the second risk category. Companies that develop and provide such applications and their clients are required to disclose specified information. Job applicants who believe that they have been discriminated against by AI hiring applications can file EU class action suits against the AI application vendor and/or the recruiting organization, and both are legally required to explain the provenance of the AI application’s outputs.

In the U.S., a Presidential Executive Order on “Safe, Secure and Trustworthy AI”<sup>16</sup> requires developers of AI applications relating to national security and national public health and safety to share their safety test results with the government. This order also directs government departments that deal with energy and homeland security to develop standards, tools and tests to help ensure that AI systems are safe, secure and trustworthy. Specific U.S. laws at the sector level have not yet been developed. In the meantime, many U.S. organizations are following broad

AI guidelines from the National Institute of Standards and Technology.<sup>17</sup>

Given such complexities in the legal environment, AI software vendors should build credibility with clients by demonstrating their knowledge of and seriousness in addressing regulatory requirements and by showing they are among the leading and proactive voices in their industries. HireVue did this by creating internal policy documents that explain how its actions address regulatory requirements and including some of the material from these documents in their communications with clients. Being alert to regulatory requirements and communicating compliance with clients also builds trust with them. These documents also provide evidence of compliance with rules existing at the time of execution in case of future litigation.

### 4. Recognize that Transparency Is an Ongoing Endeavor

AI software vendors should recognize that transparency is a set of continual processes for observing, understanding and communicating about key activities that shape and steer the AI application’s development and deployment. AI applications exist in a complex and changing world of technological advancement and societal change. As technology evolves, so does the ability of AI application developers to explain AI algorithms. For example, advances in the capabilities of large language models can be leveraged to develop explanations of why HireVue’s hiring application gives certain interview responses a good score for a specific competency.

AI software vendors often see transparency as a one-time response to a specific request or requirement (e.g., disclosing information about the AI algorithm because an affected person filed a lawsuit or to comply with a legal requirement). We recommend, however, that transparency should be an ongoing endeavor. Consider, for example, concerns about algorithmic bias resulting in the underrepresentation of women or certain ethnic groups in hiring. In this example, the actions relating to internal transparency practices should focus on continual calibration

15 Lebovitz, S., Lifshitz-Assaf, H. and Levina, N. “To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis,” *Organization Science* (33:1), January 2022, pp. 126-148.

16 *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*, October 30, 2023, available at <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

17 *A Plan for Global Engagement on AI Standards*, National Institute of Standards and Technology, July 26, 2024, available at <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-5.pdf>.

and potential correction and dissemination of the results among different functions associated with the application’s development and deployment (e.g. IT function, marketing/sales function, product development/R&D function). Domain experts and data scientists should keep client-facing functions informed about potential changes in the AI application based on such analyses and on their own research.

Transparency practices must therefore be open to change, mirroring ongoing bias analysis of the AI algorithm, emerging computational and statistical modeling methods and new regulations. Thus, providing AI transparency is never finished—complete and once-and-for-all transparency cannot be achieved. Rather, by making the underlying processes visible, transparency empowers organizations that develop and implement AI applications to continually and consciously understand, challenge and re-create their own practices of developing AI models and make their consequences clearer to internal departments and clients as well as industry and regulatory bodies.

## 5. Involve All Organizational Functions in AI Transparency Practices

The five types of transparency practices reveal how several different stakeholders across the consecutive steps of AI development and deployment shape and are affected by an AI application. Multiple departments and their values and activities play vital roles in shaping AI transparency. Given the lack of a unifying view of transparency across HireVue’s different departments, the company’s psychologists, domain experts, sales/marketing executives and IT leaders worked together to go beyond just algorithmic transparency.

This approach is particularly relevant in organizations developing AI applications for clients or for their own use or working with AI application vendors to deploy specific AI solutions. Departments involved with an AI application typically include the data science and IT functions (which work directly on the application), those relating to the application’s domain (industrial psychology in the case of HireVue), client-facing functions such as marketing/sales and functions interfacing with

external entities such as regulatory bodies. We recommend that individuals from these departments engage in transparency-related actions during each step of the AI application’s journey, starting with client engagement for a new AI application and going all the way to the post-launch assessment of that application. These actions will provide clarity on model assumptions and model traceability, help clients comprehend the AI model’s components and data, and provide explanations of the application’s benefits, limitations and (un)expected consequences of use.

## Concluding Comments

As an early mover, HireVue attracted attention for its use of AI in hiring. Various stakeholders—applicants, recruiters, hiring managers, journalists and academics—initially questioned the technology because of understandable concerns that it could do harm, either by failing to provide a valid and accurate assessment of an applicant’s potential in the job role or by increasing bias toward or against certain groups in the hiring process.<sup>18</sup> Because of these concerns, there was both strategic and ethical value in HireVue’s application being transparent.

To demonstrate to both external and internal stakeholders that much of the concern about its AI hiring application was unwarranted or had been addressed through careful development procedures, the need for HireVue to establish transparency was strategically important. Transparency would provide an accurate picture of the effects of using AI in the hiring process. The ethical importance of transparency stemmed from employment being a core element of economic and social well-being and the potential role of the AI hiring application to improve equitable access to economic opportunities, especially for historically underrepresented groups.

The HireVue case described in this article provides the AI developer and implementation communities with an example of the transparency practices used to document application outcomes in a sensitive application domain. Our

18 Maurer, R. *HireVue Discontinues Facial Analysis Screening*, SHRM, February 3, 2021, available at <https://www.shrm.org/topics-tools/news/talent-acquisition/hirevue-discontinues-facial-analysis-screening>.

recommendations for creating an AI “glass box” derived from this case shift transparency from an opaque technical AI component (a black box) to the larger sociotechnical system in which human institutions, processes and operators interact with the technical components to provide a systemic, measurable and interpretable result.

The HireVue case also shows how examining issues such as bias in one stage of a complex process such as hiring can lead to unintended beneficial consequences of illuminating problems in other non-AI-supported stages. Though the design decisions made by HireVue, such as banding applicants into three categories, may not be relevant to all AI applications, the case indicates the level of detail necessary for preparing these applications and their documentation.

## Appendix: Research Method

To study how HireVue implemented transparency for its AI hiring application, we conducted semi-structured interviews with HireVue staff involved in developing and deploying the application. The interviews were carried out by the first two authors over a 10-month period in 2021 and 2022. The interview with HireVue’s CEO identified the organization’s different departments, and we subsequently interviewed the heads of the occupational psychology, data science and assessment innovations departments. These participants suggested additional interviewees in their own and other departments.<sup>19</sup>

Our goal was to understand the different transparency-related actions throughout the lifecycle of HireVue’s AI application. In total, we conducted 28 interviews with 21 members of the organization (11 women and 10 men). The table below lists the interviewees, the number of interviews and their departments and functions. Eight interviewees had a Ph.D. in a field related to their job role, such as psychology or data science. Six with client-facing roles were interviewed to gain insights into how the AI application was designed and implemented. The interview protocol, described in the second table below, covered the stages of development

and deployment of the AI application, the organization’s approaches to transparency and establishing trust with clients, and the limitations of these practices.

The interviews, averaging 45-60 minutes, were conducted via video on Microsoft Teams. Participants signed a nondisclosure agreement with the research team, enabling them to share their screens and reveal documents such as explainability forms and and/or reports with explanations accompanying AI scores, usually exclusive to clients.

Data collection and analysis of the emerging themes were carried out concurrently with the interviews, with the first two authors meeting regularly to discuss the results. We used thematic analysis to analyze the interviews using NVivo11.<sup>20</sup> We also gathered data from HireVue’s publications, client materials and external audit reports. Specifically, we read seven white papers that described how the AI application was designed and implemented and provided information on its lifecycle and 13 peer-reviewed academic publications written by HireVue’s occupational psychologists and data scientists that explained how the organization enhanced its knowledge of AI assessment and how it disseminated its transparency strategy.

We also read two external audit reports on HireVue’s AI hiring application and occupational psychology approaches carried out by ORCAA and Landers Workforce Science. These audits covered things such as adherence to known best practices in occupational psychology and described the steps the organization took to ensure adherence. This helped us understand HireVue’s approaches to creating credibility. We also accessed client material/infographics that visualized and explained how the organization developed the assessment and its transparency actions in the design and development processes.

We conducted three rounds of thematic coding, stopping when no new codes emerged (i.e., when we reached thematic saturation).<sup>21</sup> In the first round, we identified the AI application’s lifecycle stages. In the second, we identified the organization’s activities with regard to

19 Cohen, N. and Arieli, T. “Field Research in Conflict Environments: Methodological Challenges and Snowball Sampling,” *Journal of Peace Research* (48:4), July 2011, pp. 423-435.

20 Boyatzis, R. E. *Transforming Qualitative Information: Thematic Analysis and Code Development*, SAGE Publications, 1998.

21 Guest, G., Bunce, A. and Johnson, L. “How Many Interviews Are Enough? An Experiment with Data Saturation and Variability,” *Field Methods* (18:1), February 2006, pp. 59-82.



## Interviewees’ Roles, Sex and Departments

Role in HireVue	Sex (Number of Interviews)	Department and Its Function
CEO	1 male (1)	<i>Senior management:</i> Strategic and corporate head.
Chief Occupational Psychologist (Ph.D.)	1 male (2)	
Director of Assessment Psychometrics, (Ph.D.)	1 male (2)	
Senior Occupational Psychology Consultants (Ph.D.)	2 females (3)	
Occupational Psychology Consultants	2 females (2)	
Head of Assessments (Ph.D.)	1 male (1)	<i>Assessment implementation:</i> Conducts comprehensive research on the AI application (e.g., monitoring the presence of representative demographic data, updating the algorithms).
Director of Assessments Innovation, (Ph.D.)	1 female (1)	
Chief Data Scientist (Ph.D.)	1 female (1)	
Senior Data Scientist (Ph.D.)	1 male (2)	<i>Data science:</i> Builds the assessment, conducts data analytics, researches diverse methods for mitigating bias and rigorously tests them with the AI models, performing ongoing maintenance of the AI hiring application.
Data Scientist	1 female (1)	
Director of Product	1 male (1)	
Product Managers	1 male (1) 1 female (2)	<i>Product development:</i> Advises on the strategic development of the AI hiring application.
Senior Vice President of Client/ Customer Success	1 male (1)	
Director of Client/ Customer Success	1 female (2)	
Principal Solutions Architects	1 male 2 female (4)	
Solution Architect	1 male (1)	<i>Client-facing roles:</i> Seeks to understand the client’s operational context and offers tailored advice on using the AI hiring application, provides ongoing consultancy for the client throughout the AI application implementation pipeline, gathers and incorporates client feedback, demonstrates the application’s value post-implementation and facilitates contract renewals with clients.

transparency and establishing trust with clients. In the third round, we mapped the codes from the second round to the first round—i.e., we mapped the organization’s transparency activities to the lifecycle stages. Finally, using an iterative process we merged the codes identified in the previous rounds to develop high-level categories, resulting in the five types of transparency practices.

Our research followed a qualitative participatory research method.<sup>22</sup> According to this method, research participants as co-authors can act as a strong reference and/or expert group

in validating the findings as well as in linking the emerging implications from the study back to their day-to-day practice.<sup>23</sup> We therefore engaged further with two participants—HireVue’s Chief Occupational Psychologist and Chief Data Scientist, both of whom are subject matter experts and have high-level insight into the HireVue’s operations. They joined us as co-authors in 2023, by which time the draft of the findings had been written up. They collaborated on the article by reviewing and editing it. In

22 Lister, K., Coughlan, T., Kenny, I., Tudor, R. and Iniesto, F. “Taylor, the Disability Disclosure Virtual Assistant: A Case Study of Participatory Research with Disabled Students,” *Education Sciences* (11:10), September 2021, pp. 587.

23 Lindheim, T. “Participant Validation: A Strategy to Strengthen the Trustworthiness of Your Study and Address Ethical Concerns,” in Espedal, G., Løvaas, B. J., Sirris, S. and Wæraas, A. (eds.) *Researching Values: Methodological Approaches for Understanding Values Work in Organisations and Leadership*, Springer International Publishing, 2022, pp. 225-239.



## Interview Protocol

Topic	Questions
General familiarity	<ol style="list-style-type: none"> <li>1. Tell us about yourself.</li> <li>2. Describe what you do at HireVue.</li> </ol>
AI hiring application	<ol style="list-style-type: none"> <li>3. Tell us about the AI hiring application that you offer.</li> <li>4. What are the stages of managing this application and configuring it for a client? What are your touch points on the application? Can you talk us through the product’s development cycles?</li> </ol>
AI hiring application and the organization’s transparency strategies	<ol style="list-style-type: none"> <li>5. In your view, what makes HireVue practices transparent? How does HireVue implement transparency in its processes?</li> <li>6. How is transparency being communicated? When you say that “we are transparent,” what would be the manifestation of that?</li> <li>7. Can you reflect on what makes your AI product transparent, what specific features, or the way it is set up, make you think that it is transparent?</li> <li>8. Do you have to get your product approved as to the U.S. minimum legal requirement? Do you have to show proof of that?</li> </ol>
Explainability vis-à-vis clients	<ol style="list-style-type: none"> <li>9. Do you walk the client through the product, do you explain the logic behind things? How do you help them understand the product?</li> <li>10. Does your team sometimes act as a translator, especially when it comes to establishing trust and explaining to the client how the models are built, what they mean?</li> <li>11. When you realize that a certain explainability strategy does not work, because the input you provide does not make sense to the client, how does this come about, and how do you solve it? Do you collaborate with other teams to tweak your approach?</li> <li>12. What do you mean by transparency when you speak to clients?</li> <li>13. How do you see trust unfolding on the part of the client based on what you do in terms of giving out information to them?</li> </ol>
Limitations of transparency practices	<ol style="list-style-type: none"> <li>14. Are there nuances in how you speak to the client, what you tell them, what you don’t tell them, and why you tell them things?</li> <li>15. In your interactions with the client, how would you describe the transparency aspect? How much information is too much, how much is too little? How do you make the balance? And how do you decide how to approach this for different clients?</li> <li>16. Does transparency come into conflict with any other ethical principles in your practice? How are those conflicts resolved in HireVue?</li> </ol>
Feedback	<ol style="list-style-type: none"> <li>17. Do your clients tell you anything about the product that triggers it to change?</li> <li>18. Do you get any kind of feedback on your products from clients? Is it qualitative feedback or data feedback?</li> <li>19. Have you ever received pushback from clients on what you are saying to them (e.g., “we don’t agree with this” or an extreme case “no, we don’t even want to go ahead with this”)?</li> <li>20. Do you review your product after it has been deployed to the client? How does that happen? Within what timeframe? Is it included in the contract with the client?</li> <li>21. What external entities do you liaise with (regulatory bodies, scientists, data companies, etc.)?</li> </ol>

particular, they helped validate our findings and enhance their contextual relevance, currency and reliability, especially the technical details and retrospective linking back of the interpretation of the data and findings to the organizational context of HireVue. They also offered further

insights into the practical implications of the identified transparency practices.

## About the Authors

### Monideepa Tarafdar

Monideepa Tarafdar (mtarafdar@umass.edu) is Charles J. Dockendorff Endowed Professor at Isenberg School of Management, University of Massachusetts Amherst, and editor-in-chief of *Journal of the Association for Information Systems*. She has received research grants of over \$1.5 million as principal investigator from, among others, Leverhulme Trust (U.K.) and the Economic and Social Science Research Council (U.K.) and has collaborated on research with organizations in the U.S., U.K., EU and India. Her work has been published in journals such as *Science*, *PNAS Nexus*, *Information Systems Research*, *Journal of the AIS*, *Journal of MIS*, *Journal of Operations Management*, *Harvard Business Review* and *Sloan Management Review*.

### Irina Rets

Irina Rets (irina.rets@open.ac.uk) is a senior research fellow at the Institute of Educational Technology, The Open University, U.K. She specializes in inclusive AI and education, focusing on how technology can promote social justice in learning and society. With over 25 publications in prestigious outlets, Irina is committed to transforming research into impact through policy-oriented outputs, developing effective impact strategies for her research projects, and creating accessible summaries of her work.

### Lindsey Zuloaga

Lindsey Zuloaga (lzuloaga@hirevue.com) is an AI thought leader and HireVue’s chief data scientist. With a doctorate in applied physics from Rice University, she combines scientific rigor with innovative thinking to advance ethical AI and algorithmic transparency. At HireVue, Lindsey spearheads initiatives that leverage AI to transform hiring practices while prioritizing fairness, accountability and inclusivity. As a speaker and advisor, she is recognized for her contributions to discussions on the responsible use of AI, shaping industry standards and fostering the ethical integration of technology in the workplace.

### Nathan Mondragon

Dr. Nathan Mondragon (nmondragon@hirevue.com) is HireVue’s chief innovation officer, responsible for creating software products driven by scientific principles. He co-led or led the creation of many science-based technology solution firsts, such as the creation and delivery of the first online selection assessment system in 1996 and helped deliver the first AI-driven pre-hire assessment solution at HireVue. Nathan has spoken with Congress and the White House on the proper methods to develop and maintain AI technology, has written over 100 manuscripts, peer-reviewed papers and book chapters and has delivered presentations and workshops on industrial and organizational psychology and technology solutions.