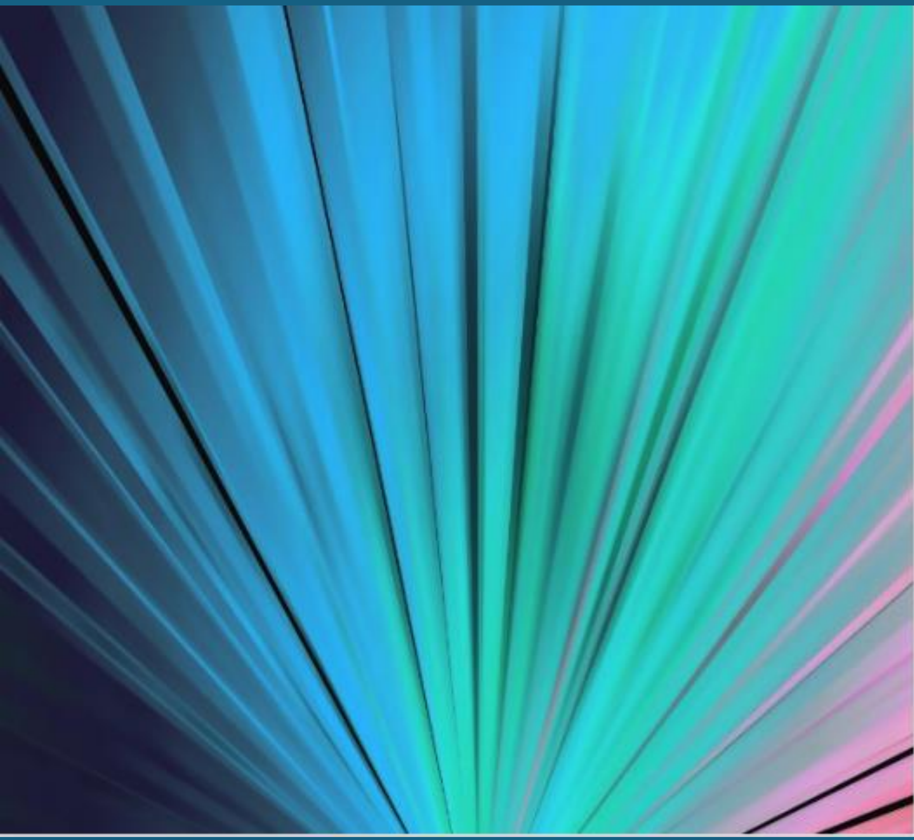

The Ultimate AI Glossary for Recruitment and HR

AI Recruitment Glossary: A Non-Technical Guide

A comprehensive guide to artificial intelligence terminology in recruitment, written for HR professionals

Executive Summary

This glossary provides HR professionals with essential AI recruitment terminology, from technical concepts to regulatory requirements. Whether you're evaluating vendors, preparing for board presentations, or implementing AI systems, this guide offers practical insights and real-world examples to help you navigate the AI recruitment landscape confidently.

Key areas covered: Technical foundations, bias and fairness, regulatory compliance, vendor evaluation, and strategic considerations for successful AI adoption.

Top 10 Terms Every HR Director Should Know

1. **Algorithmic Bias** - Risk of unfair discrimination in AI systems
2. **Hallucination** - When AI generates plausible but incorrect information
3. **Human-in-the-Loop** - Maintaining human control over AI decisions
4. **Explainable AI** - Understanding how AI makes recommendations
5. **GDPR Compliance** - Legal requirements for data protection
6. **Prompt Engineering** - Crafting effective AI instructions
7. **Predictive Analytics** - Forecasting recruitment outcomes
8. **AI Governance** - Framework for responsible AI use
9. **Model Drift** - Performance degradation over time
10. **Audit Trail** - Tracking AI decision-making processes

A

A/B Testing: A systematic method of comparing two different versions of AI recruitment processes, system configurations, or candidate experiences to determine which performs better. This involves running controlled experiments where candidates are randomly assigned to different approaches and outcomes are measured.

Real-world example: You might test two different AI-generated job descriptions for the same role, sending version A to half your target audience and version B to the other half, then measuring which generates higher-quality applications. Or you could test different AI screening criteria to see which produces better hiring outcomes.

Why it matters: A/B testing enables data-driven optimisation of your AI recruitment systems, helping you make evidence-based decisions about system configurations and candidate experiences rather than relying on assumptions.

What to look for: Systems that support controlled testing, clear statistical significance measures, and easy implementation of winning variations.

Algorithmic Accountability: The legal and ethical responsibility for AI system decisions, including clear ownership of outcomes, transparent decision-making processes, and mechanisms for addressing errors or bias. This encompasses both technical accountability (how systems work) and organisational accountability (who is responsible).

Real-world example: If an AI system unfairly rejects qualified candidates from certain backgrounds, algorithmic accountability requires clear processes for identifying the problem, determining responsibility, correcting the system, and providing remedies to affected candidates.

Why it matters: Algorithmic accountability is increasingly required by law and helps build trust with candidates, regulators, and internal stakeholders. It also helps organisations manage risks associated with AI deployment.

What to look for: Clear governance structures, defined roles and responsibilities, audit capabilities, and processes for addressing algorithmic errors or bias.

Algorithmic Bias: When an AI system consistently treats certain groups of people unfairly, often reflecting historical discrimination patterns in the data used to train it. This isn't intentional—it happens because AI systems learn from past human decisions that contained unconscious biases.

Real-world example: If your company historically hired more men for engineering roles, the AI might learn that being male is a "positive" indicator for engineering success. It would then unfairly favour male candidates, even though gender has no bearing on technical ability.

Why it matters: Algorithmic bias can expose your organisation to discrimination claims and prevent you from accessing diverse talent pools. It's also illegal under UK equality legislation.

What to look for: Regular bias testing, diverse training data, and clear processes for identifying and correcting unfair patterns in AI recommendations.

API (Application Programming Interface): A set of rules and protocols that allow different software applications to communicate and share information automatically. Think of it as a digital translator that helps your AI recruitment system talk to your existing HR software.

Real-world example: When a candidate applies through your careers page, an API might automatically transfer their details from the application system to your AI screening tool, then to your ATS, without any manual data entry.

Why it matters: Good API integration means your AI system can work seamlessly with your existing tools, reducing manual work and preventing data silos.

What to look for: Pre-built connections to popular HR systems, clear documentation, and support for custom integrations if needed.

AI Auditing / Audit Trail: The comprehensive process and capability to trace, document, and review how AI systems make decisions. This includes maintaining detailed records of data inputs, model versions, decision logic, and outcomes for compliance and accountability purposes.

Real-world example: An audit trail might show that Candidate A was scored 8.2/10 on 15th March 2024 using Model Version 2.3, based on specific data inputs (skills match 85%, experience 90%, cultural fit 75%), with the decision reviewed by Recruiter X and approved by Hiring Manager Y.

Why it matters: Audit trails are essential for GDPR compliance, employment tribunal defence, and EU AI Act requirements. They provide evidence of fair decision-making and enable investigation of bias complaints.

What to look for: Complete decision histories, version control for models, immutable logging systems, and ability to recreate historical decisions for legal defence.

AI Governance: A comprehensive framework of policies, processes, oversight mechanisms, and accountability structures for responsible AI use in hiring. This includes ethics guidelines, risk management, compliance monitoring, and strategic decision-making about AI adoption.

Real-world example: An AI governance framework might include an AI Ethics Committee with representatives from HR, Legal, IT, and business units, regular bias audits, defined approval processes for new AI tools, and clear escalation procedures for AI-related issues.

Why it matters: Strong AI governance reduces legal risks, ensures ethical AI use, builds stakeholder confidence, and provides clear accountability structures for AI-related decisions.

What to look for: Clear governance structures, regular review processes, stakeholder involvement, and integration with broader corporate governance frameworks.

AI Readiness Assessment A systematic evaluation process to determine whether an organisation is prepared to successfully adopt AI recruitment tools. This assesses data quality, technical infrastructure, legal preparedness, cultural readiness, and change management capabilities.

Real-world example: An AI readiness assessment might evaluate whether you have sufficient historical hiring data (minimum 1,000 candidate records), clean data formats, adequate IT infrastructure, legal framework understanding, and recruiter buy-in for AI adoption.

Why it matters: Proper readiness assessment prevents failed implementations, identifies preparation needs, and ensures realistic expectations for AI adoption timelines and outcomes.

What to look for: Comprehensive assessment tools, clear readiness criteria, gap analysis capabilities, and implementation roadmaps based on assessment results.

Artificial Intelligence (AI): Computer systems designed to perform tasks that typically require human intelligence, such as understanding language, recognising patterns, making decisions, or solving problems. In recruitment, AI can read CVs, match candidates to jobs, conduct initial screenings, and predict hiring success.

Real-world example: An AI system might read 1,000 CVs in minutes, identify the top 50 candidates based on job requirements, draft personalised rejection emails for unsuccessful applicants, and schedule interviews for the shortlisted candidates.

Why it matters: AI can dramatically reduce time-to-hire, improve candidate quality, and free up recruiters to focus on relationship-building and strategic work.

What to look for: Systems that enhance rather than replace human judgement, with clear explanations for their recommendations.

Attention Mechanisms: A technical approach that enables AI systems to focus on the most relevant parts of candidate profiles or job descriptions when making matching decisions, rather than treating all information equally. This mimics how humans naturally focus on key information when evaluating candidates.

Real-world example: When analysing a candidate's CV, an attention mechanism might focus heavily on recent relevant experience and key technical skills whilst giving less weight to older qualifications or less relevant roles, improving the accuracy of candidate-job matching.

Why it matters: Attention mechanisms enable more nuanced and accurate candidate evaluation by helping AI systems identify what information is most important for specific roles and contexts.

What to look for: Systems that can explain which factors they focused on for specific recommendations, demonstrated improvements in matching accuracy, and ability to adapt attention based on different role requirements.

Automated Decision-Making: When computer systems make decisions that affect people without human involvement. Under UK and EU law, individuals have the right to know when automated decisions affect them and can request human review.

Real-world example: An AI system automatically rejecting candidates who don't meet minimum qualifications, or automatically inviting certain candidates to interview based on their scores.

Why it matters: You have legal obligations to inform candidates about automated decision-making and provide human review options. Failure to comply can result in legal challenges and regulatory fines. *What to look for:* Clear processes for human oversight, candidate notification systems, and appeal mechanisms.

B

Benchmarking: The process of comparing your AI recruitment system's performance against industry standards, competitor performance, or best-in-class implementations to identify areas for improvement and validate system effectiveness.

Real-world example: You might benchmark your AI system's time-to-hire (currently 28 days) against industry average (35 days), or compare your candidate quality scores and diversity metrics against similar organisations in your sector.

Why it matters: Benchmarking helps demonstrate ROI, identify improvement opportunities, and ensure your AI systems are delivering competitive advantage rather than just meeting minimum requirements.

What to look for: Access to industry benchmark data, standardised metrics for comparison, and regular benchmarking reports that track performance over time.

BERT (Bidirectional Encoder Representations from Transformers): A foundational AI model that revolutionised natural language understanding by reading text in both directions simultaneously (forwards and backwards). BERT enables more sophisticated resume parsing, job description analysis, and candidate-job matching by understanding context and nuance in language rather than just matching keywords.

Real-world example: BERT might understand that "managed a team during digital transformation" indicates both leadership skills and change management experience, recognising the contextual relationship between these concepts. It could also understand that "spearheaded customer success initiatives" is relevant for "client relationship management" roles, even though the exact terms don't match.

Why it matters: BERT-based systems provide more accurate candidate matching, better understand implicit skills and experiences, and can identify relevant candidates who might be missed by traditional keyword searches. This leads to higher-quality shortlists and reduced time spent reviewing unsuitable applications.

What to look for: Systems that use BERT or similar transformer architectures, demonstrated improvements in matching accuracy over keyword-based systems, and ability to understand contextual relationships in job requirements and candidate profiles.

Bias Mitigation: The various techniques and strategies used to reduce unfair treatment in AI systems. This involves both technical solutions (adjusting algorithms) and process solutions (diverse training data, regular testing).

Real-world example: A vendor might use "adversarial debiasing" (a technical approach that actively counters bias during training) combined with regular testing to ensure their system doesn't favour certain demographics. They might also remove potentially biased data points like postcodes that correlate with socioeconomic status.

Why it matters: Without proper bias mitigation, AI systems can perpetuate or amplify existing inequalities, leading to legal risks and missed talent opportunities.

What to look for: Multiple bias mitigation strategies, regular third-party audits, and transparent reporting on fairness metrics.

Boolean Search: A search method using logical operators (AND, OR, NOT) to create precise queries. While traditional in recruitment, understanding Boolean logic helps you better communicate with AI systems and validate their search capabilities.

Real-world example: Searching for "(Marketing OR Digital Marketing) AND London NOT (Intern OR Graduate)" to find experienced marketing professionals in London, excluding entry-level positions.

Why it matters: Understanding Boolean logic helps you create better job descriptions and search criteria that AI systems can interpret accurately.

What to look for: AI systems that can understand complex search logic and translate natural language queries into precise searches.

C

Calibration: How well an AI system's confidence predictions match real-world outcomes. A well-calibrated system means that when it says it's 80% confident about a candidate match, it should be correct about 80% of the time.

Real-world example: If an AI system rates 100 candidates as "90% likely to succeed" in a role, around 90 of them should actually perform well if hired. If only 60 perform well, the system is overconfident and poorly calibrated.

Why it matters: Poor calibration can lead to overconfidence in AI recommendations, resulting in bad hiring decisions. Well-calibrated systems help you make more informed decisions about which candidates to prioritise.

What to look for: Evidence of calibration testing, confidence scores that match actual outcomes, and regular recalibration processes.

Candidate Experience: The overall impression and feeling job seekers have throughout your recruitment process, from initial application to final decision. AI should enhance this experience by making it faster, more transparent, and more personalised.

Real-world example: AI might provide instant acknowledgement of applications, send personalised updates on application status, offer helpful feedback on unsuccessful applications, and schedule interviews at convenient times based on candidate preferences.

Why it matters: Poor candidate experience damages your employer brand, reduces application rates, and can lead to negative reviews on platforms like Glassdoor. Good candidate experience, enhanced by AI, can become a competitive advantage.

What to look for: AI systems that prioritise candidate communication, provide transparent processes, and offer personalised interactions rather than generic automated responses.

Candidate Journey Mapping: Using AI to track, analyse, and optimise the complete candidate experience from initial awareness through to onboarding, identifying pain points and opportunities for improvement at each stage.

Real-world example: AI might track that candidates typically drop out after the initial application but before completing assessments, then recommend optimising the assessment process or providing better communication about next steps.

Why it matters: Candidate journey mapping helps improve candidate experience, reduces drop-out rates, and ensures your recruitment process is competitive and candidate-friendly.

What to look for: Systems that can track candidate interactions across multiple touchpoints, identify common drop-off points, and provide actionable recommendations for process improvements.

Candidate Matching: The process of comparing candidate profiles against job requirements to identify potential fits. AI systems can consider hundreds of factors simultaneously and identify patterns that humans might miss.

Real-world example: Instead of just matching keywords, AI might recognise that a candidate who led "digital transformation projects" has relevant experience for a "change management" role, even though the exact terms don't match. It might also weight recent experience more heavily than older qualifications.

Why it matters: Better matching leads to higher-quality shortlists, reduced time-to-hire, and improved hiring success rates. It also helps identify candidates who might be overlooked by traditional keyword matching.

What to look for: Systems that go beyond keyword matching, consider career progression patterns, and can explain why candidates are recommended.

Candidate Net Promoter Score (NPS): A metric measuring how likely candidates are to recommend your organisation's recruitment process to others, providing insights into candidate satisfaction with AI-enhanced recruitment experiences.

Real-world example: After completing your AI-powered recruitment process, candidates might be asked "How likely are you to recommend our application process to a friend?" with responses used to calculate an NPS score and identify areas where AI is enhancing or detracting from candidate experience.

Why it matters: Candidate NPS helps measure the impact of AI on employer brand, identifies areas for improvement in AI-human interactions, and provides early warning of candidate experience issues.

What to look for: Automated NPS collection systems, benchmarking against industry standards, and ability to correlate NPS scores with specific AI system features or interactions.

Chatbot: An AI-powered conversational interface that can interact with candidates through text or voice, answering questions, collecting information, and guiding them through processes. Modern chatbots can handle complex conversations and escalate to humans when needed.

Real-world example: A chatbot might greet candidates on your careers page, answer questions about company culture, help them find suitable roles, collect initial screening information, and schedule interviews—all while maintaining a friendly, professional tone.

Why it matters: Chatbots provide 24/7 candidate support, reduce response times, and free up recruiters for more strategic work. They also provide consistent information and can handle multiple conversations simultaneously.

What to look for: Natural conversational abilities, seamless handoff to humans, multilingual support, and integration with your existing systems.

Clustering: An AI technique that automatically groups similar candidates, roles, or hiring patterns without being explicitly told what to look for, revealing hidden segments and relationships in your recruitment data.

Real-world example: Clustering might reveal that your successful sales hires naturally group into three distinct profiles: "relationship builders" with strong interpersonal skills, "hunters" with aggressive target-hitting records, and "consultative sellers" with deep technical knowledge, helping you diversify your sourcing strategies.

Why it matters: Clustering helps identify hidden patterns in successful hires, reveals new candidate segments to target, and can improve matching by understanding different "types" of suitable candidates for each role.

What to look for: Systems that can identify meaningful candidate clusters, explain cluster characteristics, and use clustering insights to improve matching and sourcing strategies.

Cold Outreach: AI-powered personalised messaging to passive candidates who haven't expressed interest in your organisation, using intelligent targeting and customised communications to generate interest and engagement.

Real-world example: AI might identify a software engineer whose skills and career trajectory suggest they'd be interested in your senior developer role, then craft a personalised message referencing their recent project work and explaining why your opportunity aligns with their career goals.

Why it matters: Cold outreach enables access to passive talent pools, reduces dependence on job boards, and can help identify candidates who wouldn't have applied through traditional channels.

What to look for: Intelligent candidate identification, personalised message generation, integration with professional networks, and tracking of outreach effectiveness.

Cold Start Problem: The challenge faced when implementing AI systems in organisations with limited historical data. New systems need examples to learn from, but new companies or departments may not have enough past hiring data to train effective AI models.

Real-world example: A startup with only 20 employees wants to implement AI recruitment but doesn't have enough historical hiring data to train a system. The AI vendor might use pre-trained models based on similar companies, external data sources, or industry benchmarks to compensate.

Why it matters: Understanding the cold start problem helps set realistic expectations for AI implementation timelines and performance. It also influences your choice of vendor and implementation strategy.

What to look for: Vendors with pre-trained models, external data enrichment capabilities, and realistic timelines for achieving optimal performance.

Concept Drift: The phenomenon where the relationships between job requirements and candidate success change over time, making older training data less relevant. This is particularly relevant in fast-changing industries or during major market shifts.

Real-world example: Pre-2020, remote work skills weren't essential for most roles. Post-pandemic, these skills became crucial. An AI system trained on pre-2020 data might not properly value candidates with strong remote collaboration abilities.

Why it matters: Without regular updates, AI systems can become increasingly inaccurate over time, leading to poor hiring decisions and missed opportunities.

What to look for: Systems with regular retraining schedules, drift detection capabilities, and processes for incorporating new market realities.

Conversational AI: A systemsthat can engage in natural, human-like conversations, understanding context, maintaining conversation flow, and providing relevant responses. In recruitment, this powers chatbots, interview assistants, and candidate communication systems.

Real-world example: A conversational AI might conduct initial candidate screenings, asking follow-up questions based on responses, clarifying ambiguous answers, and adapting its conversation style to match the candidate's communication preferences.

Why it matters: Conversational AI can provide more engaging candidate experiences, gather better information through natural dialogue, and handle complex queries that simple chatbots cannot.

What to look for: Natural conversation flow, context understanding, ability to handle complex queries, and appropriate escalation to human recruiters.

D

Data Portability: The legal requirement under GDPR and other privacy regulations for candidates to be able to transfer their personal data between different systems, including AI recruitment platforms.

Real-world example: A candidate might request to transfer their profile data, assessment results, and application history from your AI recruitment system to another employer's system, requiring you to provide the data in a structured, commonly used format.

Why it matters: Data portability is a legal requirement that affects how AI systems store and manage candidate data, and influences system design and vendor selection decisions.

What to look for: Systems that support standard data export formats, clear data portability procedures, and compliance with relevant privacy regulations.

Data Quality: The accuracy, completeness, consistency, and relevance of the information used to train and operate AI systems. Poor data quality leads to poor AI performance—the classic "garbage in, garbage out" principle.

Real-world example: If your historical hiring data contains incomplete candidate profiles, outdated job descriptions, or inconsistent performance ratings, the AI system will learn from these flaws and make poor recommendations. Good data quality means complete profiles, standardised job descriptions, and consistent performance metrics.

Why it matters: Data quality directly impacts AI effectiveness. Poor data can lead to biased recommendations, inaccurate matching, and compliance issues.

What to look for: Data cleansing services, quality assessment tools, and ongoing data maintenance processes.

Deep Learning: A type of AI that uses neural networks with multiple layers to learn complex patterns from large amounts of data. It's particularly effective at understanding unstructured data like text, images, and speech.

Real-world example: Deep learning might analyse thousands of successful employee profiles to identify subtle patterns that predict job success—perhaps discovering that candidates who use certain phrases in their cover letters or have specific combinations of experiences tend to perform better.

Why it matters: Deep learning can identify complex patterns that traditional methods miss, potentially leading to better candidate insights and more accurate predictions.

What to look for: Systems that can handle complex, unstructured data and provide insights that go beyond simple rule-based matching.

Demographic Parity: A fairness metric requiring that AI decisions are made at the same rate across different demographic groups. For example, the system should shortlist similar percentages of male and female candidates.

Real-world example: If 30% of male applicants are shortlisted, then approximately 30% of female applicants should also be shortlisted. Significant deviations might indicate bias.

Why it matters: Demographic parity helps ensure fair treatment across groups and can help identify potential bias in AI systems.

What to look for: Regular monitoring of selection rates across different groups, automated alerts for significant disparities, and clear processes for addressing imbalances.

Data Poisoning: When malicious actors deliberately insert misleading, biased, or corrupted data into training datasets to sabotage AI system performance or introduce specific vulnerabilities. This is an increasing concern with open-source or user-contributed models in HR tech.

Real-world example: An attacker might submit fake candidate profiles with misleading information designed to skew AI recommendations or inject biased data that causes the system to discriminate against certain groups.

Why it matters: Data poisoning can compromise AI system reliability, introduce bias, and potentially expose organisations to legal and reputational risks from flawed hiring decisions.

What to look for: Vendors with data validation processes, secure data collection methods, anomaly detection systems, and clear data provenance tracking.

Diversity and Inclusion (D&I): Efforts to ensure fair representation and treatment of all groups in recruitment and employment. AI can support D&I goals by reducing unconscious bias, but it requires careful monitoring to avoid perpetuating existing inequalities.

Real-world example: An AI system might help achieve D&I goals by anonymising candidate information during initial screening, expanding sourcing to underrepresented talent pools, and flagging potential bias in job descriptions. However, it could also perpetuate bias if trained on historical data that reflected past discrimination.

Why it matters: D&I is both a legal requirement and a business imperative. Diverse teams perform better, and inclusive hiring practices help access wider talent pools.

What to look for: AI systems designed with D&I in mind, regular bias testing, and features that support inclusive hiring practices.

E

Embedding Models: The mathematical representation of text, skills, experiences, and job requirements as numerical vectors in multi-dimensional space. These models enable AI systems to understand semantic relationships and similarities between different concepts, even when they use different words.

Real-world example: An embedding model might represent "project management," "programme coordination," and "initiative leadership" as similar vectors because they share semantic meaning, allowing the AI to recognise that candidates with any of these experiences might be suitable for project management roles.

Why it matters: Embedding models are fundamental to modern AI recruitment systems, enabling semantic matching that goes far beyond keyword searches. They allow systems to understand that "Python developer" and "software engineer with Python experience" represent similar concepts.

What to look for: Systems that use modern embedding models, can demonstrate semantic understanding beyond keyword matching, and regularly update their embeddings to reflect evolving language and job market changes.

Ensemble Methods: Advanced AI approaches that combine multiple different models or algorithms to improve prediction accuracy and reduce bias, rather than relying on a single AI system.

Real-world example: An ensemble system might combine a neural network that's good at analysing CVs, a natural language processing model that assesses cultural fit from cover letters, and a predictive model that forecasts performance, creating more robust and accurate candidate evaluations.

Why it matters: Ensemble methods typically provide more accurate predictions than single models, reduce the risk of systematic bias, and create more robust AI systems that perform well across diverse scenarios.

What to look for: Systems that combine multiple AI approaches, demonstrated improvements in accuracy over single-model systems, and transparent explanations of how different models contribute to final recommendations.

Equalised Odds: A fairness metric ensuring that the AI system's accuracy is consistent across different demographic groups. The system should be equally good at identifying strong candidates regardless of their background.

Real-world example: If the AI correctly identifies 80% of high-performing white candidates, it should also correctly identify around 80% of high-performing candidates from ethnic minorities. Significant differences might indicate bias.

Why it matters: Equalised odds helps ensure that AI systems don't systematically disadvantage any group, maintaining fairness in recruitment processes.

What to look for: Regular accuracy testing across different groups, monitoring of performance disparities, and processes for addressing identified issues.

EU AI Act: Comprehensive European Union regulation governing artificial intelligence systems, which classifies AI used in recruitment and employee management as "high-risk" systems requiring strict compliance measures, transparency requirements, and human oversight.
Real-world example: Under the EU AI Act, recruitment AI systems must undergo conformity assessments, maintain detailed documentation, provide transparency to candidates about AI use, and ensure human oversight of significant decisions.

Why it matters: Non-compliance with the EU AI Act can result in fines up to €35 million or 7% of global turnover. UK companies using AI for EU candidates or employees must also comply.

What to look for: Vendors preparing for EU AI Act compliance, conformity assessment processes, and systems designed to meet high-risk AI requirements.

Explainable AI: AI systems that can provide clear, understandable explanations for their decisions and recommendations. Instead of just providing a recommendation, these systems explain their reasoning in terms that humans can understand and evaluate.

Real-world example: Rather than just saying "Sarah is a 85% match for this role," an explainable AI might say "Sarah scores highly because of her 5 years of relevant experience (30% of score), strong technical skills in required areas (25%), leadership experience (20%), and cultural fit indicators (10%). Areas for consideration include limited experience with our specific industry (reducing score by 5%)."

Why it matters: Explainability is crucial for legal compliance, building trust with recruiters, and enabling informed decision-making. It also helps identify potential bias or errors in AI reasoning.

What to look for: Clear, specific explanations for recommendations, ability to drill down into reasoning, and explanations tailored to different audiences (recruiters vs. legal teams).

External Data Sources: Information from outside your organisation that AI systems might use to enrich candidate profiles or improve matching accuracy. This could include professional networks, skills databases, market intelligence, or public records.

Real-world example: An AI system might supplement a candidate's CV with information from their LinkedIn profile, GitHub contributions, professional certifications, or industry awards. It might also use market data to understand current salary expectations or skills demand.

Why it matters: External data can provide more complete candidate pictures and help identify qualified candidates who might not have detailed CVs. However, it raises privacy and accuracy concerns.

What to look for: Transparent data sourcing, consent management, data accuracy verification, and compliance with privacy regulations.

F

Fairness Metrics: Mathematical measures used to assess whether an AI system treats different groups fairly. There are multiple ways to define fairness, and different metrics can sometimes conflict with each other.

Real-world example: You might measure fairness by comparing selection rates across groups (demographic parity), accuracy across groups (equalised odds), or ensuring that similar candidates receive similar scores regardless of group membership (individual fairness).

Why it matters: Different fairness metrics can give different results, so it's important to understand which metrics matter most for your organisation and legal requirements.

What to look for: Multiple fairness metrics, regular monitoring, and clear processes for addressing fairness issues when they arise.

False Positive: When the AI incorrectly identifies a candidate as suitable when they're actually not a good fit. This leads to wasted time interviewing poorly matched candidates and potentially bad hires.

Real-world example: An AI system might rate a candidate highly because they have impressive qualifications on paper, but miss that their actual experience doesn't align with the role requirements, or that they lack essential soft skills.

Why it matters: High false positive rates waste recruitment resources and can lead to poor hiring decisions. However, some false positives are inevitable—the goal is to find the right balance.

What to look for: Systems with good precision rates, ability to validate predictions against actual outcomes, and processes for learning from false positives.

False Negative: When the AI incorrectly rejects a candidate who would actually be suitable for the role. This is particularly concerning because it means missing out on potentially excellent hires.

Real-world example: An AI system might reject a candidate because they don't have a traditional educational background, even though they have extensive practical experience and strong skills that make them ideal for the role.

Why it matters: False negatives represent missed opportunities and can perpetuate bias if certain groups are systematically undervalued by the AI system.

What to look for: Systems with good recall rates, regular validation against actual hiring outcomes, and processes for identifying and correcting systematic errors.

Fairness Through Unawareness: An outdated and ineffective strategy that attempts to avoid bias by excluding protected characteristics (age, gender, ethnicity) from AI models. This approach often fails because proxy variables can still reveal protected information.

Real-world example: Removing gender information from candidate profiles while retaining postcode data, university information, or previous job titles—all of which can correlate with gender and enable indirect discrimination.

Why it's problematic: Fairness through unawareness provides false security while allowing discrimination through proxy variables. It's now recognised as insufficient for ensuring fair AI systems.

What to look for: Vendors who understand the limitations of this approach and use more sophisticated bias mitigation techniques like adversarial debiasing or fairness constraints.

Feedback Loop: The process by which AI systems learn from their previous decisions and outcomes to improve future performance. When managed properly, this leads to continuous improvement. When managed poorly, it can amplify existing biases.

Real-world example: An AI system learns that candidates from certain universities tend to perform well in your organisation. If this pattern is based on genuine job-relevant factors, the feedback loop improves accuracy. If it's based on unconscious bias in performance evaluations, the loop amplifies discrimination.

Why it matters: Feedback loops are essential for AI improvement but require careful management to avoid amplifying bias or creating self-reinforcing patterns.

What to look for: Carefully designed feedback mechanisms, bias monitoring in feedback loops, and processes for breaking negative feedback cycles.

Fine-tuning: The process of adapting pre-trained AI models to work specifically for your organisation's recruitment needs, using your historical hiring data and outcomes to improve accuracy and relevance.

Real-world example: A general AI model might be fine-tuned using your company's successful hiring data to better understand what "cultural fit" means in your specific context, or to recognise that certain combinations of skills and experience predict success in your particular industry.

Why it matters: Fine-tuning enables AI systems to adapt to your organisation's unique requirements, culture, and success patterns, leading to more accurate recommendations and better hiring outcomes than generic, one-size-fits-all systems.

What to look for: Vendors that offer fine-tuning services, clear processes for incorporating your data, and demonstrated improvements in performance after fine-tuning for similar organisations.

G

GDPR (General Data Protection Regulation): European privacy legislation that governs how personal data must be collected, processed, stored, and protected. It gives individuals significant rights over their data and imposes strict obligations on organisations.

Real-world example: Under GDPR, candidates have the right to know what data you're collecting, how you're using it, and can request copies of their data or ask for it to be deleted. If your AI system processes candidate data, it must comply with these requirements.

Why it matters: GDPR non-compliance can result in fines up to 4% of global turnover. AI systems that process personal data must be designed with privacy protection built in.

What to look for: GDPR-compliant data processing, clear privacy controls, data minimisation principles, and robust consent management.

Generative AI: AI systems that can create new content, such as job descriptions, interview questions, candidate communications, or assessment questions. Popular examples include ChatGPT, GPT-4, and similar large language models.

Real-world example: Generative AI might write personalised rejection emails that provide constructive feedback, create job descriptions optimised for different demographics, or generate interview questions tailored to specific roles and candidate backgrounds.

Why it matters: Generative AI can dramatically reduce content creation time and help personalise candidate communications at scale. However, it requires careful monitoring for accuracy and appropriateness.

What to look for: Quality control processes, ability to customise output style and tone, and integration with existing recruitment workflows.

H

Hallucination: When AI systems generate outputs that are plausible sounding but factually incorrect, irrelevant, or fabricated. This is particularly problematic with generative AI used for creating job descriptions, candidate communications, or interview questions.

Real-world example: An AI system might generate a job description mentioning non-existent software tools, create candidate emails with incorrect company information, or suggest interview questions based on outdated industry practices.

Why it matters: Hallucinations can damage employer brand, confuse candidates, and lead to poor hiring decisions. They're especially concerning when AI-generated content isn't properly reviewed.

What to look for: Systems with hallucination detection, human review processes for AI-generated content, and clear warnings about the need to verify AI outputs.

Human-in-the-Loop: AI systems designed to work alongside humans rather than replace them. Critical decisions always involve human judgement, with AI providing support, recommendations, and efficiency improvements.

Real-world example: An AI system might screen applications and provide ranked recommendations with explanations, but the final decision to interview or hire candidates remains with human recruiters and hiring managers.

Why it matters: Human-in-the-loop systems help maintain accountability, provide better candidate experience, and ensure that human judgement and values remain central to hiring decisions.

What to look for: Clear human oversight points, easy override capabilities, and systems that enhance rather than replace human decision-making.

Human Oversight: The requirement for meaningful human involvement in AI-assisted decision-making processes. This is particularly important for decisions that significantly affect people's lives, such as hiring.

Real-world example: While AI might automatically screen applications and rank candidates, human recruiters must review the recommendations, consider additional factors, and make final decisions about interviews and offers.

I

Integration: The process of connecting AI recruitment systems with your existing HR technology stack, ensuring seamless data flow and coordinated workflows across different platforms.

Real-world example: Your AI system might automatically pull job requirements from your HRIS, receive applications from your careers page, share candidate data with your ATS, and sync interview feedback with your performance management system.

Why it matters: Good integration eliminates data silos, reduces manual work, and ensures consistent candidate information across all systems.

What to look for: Pre-built connectors to popular HR systems, robust APIs, comprehensive documentation, and support for custom integrations.

Intersectionality: The recognition that people may belong to multiple identity groups simultaneously (e.g., being both female and from an ethnic minority) and may face compound discrimination that differs from the sum of individual discriminations.

Real-world example: An AI system might show no bias against women individually or ethnic minorities individually, but still discriminate against women from ethnic minorities due to the intersection of these characteristics.

Why it matters: Simple bias testing that looks at groups individually might miss intersectional discrimination, leading to continued unfair treatment of multiply marginalised candidates.

What to look for: Bias testing that considers intersectional groups, not just individual characteristics, and systems designed to detect complex discrimination patterns.

J

Job Matching: The comprehensive process of comparing candidate profiles against job requirements to identify potential fits. Modern AI systems can consider hundreds of factors simultaneously and identify complex patterns that indicate compatibility.

Real-world example: Beyond matching keywords, AI might recognise that a candidate's experience leading "customer success initiatives" is relevant for a "client relationship management" role, consider their career progression pattern, evaluate cultural fit indicators, and assess their potential for growth in the position.

Why it matters: Sophisticated job matching leads to higher-quality shortlists, better hiring outcomes, and more efficient recruitment processes.

What to look for: Systems that go beyond simple keyword matching, consider context and career patterns, and can explain their matching logic.

Job Parsing: The AI's ability to break down job descriptions into component parts such as required skills, qualifications, experience levels, and responsibilities. This structured understanding enables better candidate matching.

Real-world example: An AI system might read a job description and identify that it requires "5+ years Python experience" (technical skill), "team leadership experience" (soft skill), "bachelor's degree in computer science" (qualification), and "experience with agile methodologies" (process knowledge).

Why it matters: Accurate job parsing is fundamental to effective matching and ensures that the AI understands what you're actually looking for in candidates.

What to look for: Systems that can handle complex job descriptions, understand implicit requirements, and structure information in useful ways.

K

Key Performance Indicators (KPIs): Measurable metrics used to evaluate the success and effectiveness of your AI recruitment system. These should align with your broader recruitment and business objectives.

Real-world example: KPIs might include time-to-hire (reduced from 45 to 30 days), candidate quality scores (increased from 7.2 to 8.1 out of 10), diversity metrics (improved representation across all levels), cost-per-hire (reduced by 25%), and candidate satisfaction scores (increased from 6.8 to 8.3).

Why it matters: Clear KPIs help demonstrate ROI, identify areas for improvement, and ensure that AI implementation delivers measurable business value.

What to look for: Systems that can track relevant KPIs, provide clear reporting, and help you benchmark against industry standards.

Knowledge Base: The comprehensive information repository that an AI system uses to understand jobs, skills, industries, and candidates. This includes structured data (skills databases) and learned patterns (from training data).

Real-world example: A knowledge base might include understanding that "Java" is a programming language, that "project management" skills are valuable across industries, that "MBA" is a postgraduate qualification, and that candidates with certain experience patterns tend to succeed in specific roles.

Why it matters: The quality and comprehensiveness of the knowledge base directly impact the AI system's ability to make accurate matches and recommendations.

What to look for: Comprehensive, regularly updated knowledge bases, industry-specific understanding, and clear processes for incorporating new information.

L

Large Language Model (LLM): AI systems trained on vast amounts of text data that can understand and generate human-like language. These models power many modern AI recruitment tools, enabling them to read CVs, write job descriptions, and communicate with candidates.

Real-world example: An LLM might read a candidate's CV and understand that their experience as a "customer success manager" involved skills relevant to "account management" roles, even if the exact terms don't match. It might also help write personalised outreach messages to passive candidates.

Why it matters: LLMs enable more sophisticated understanding of language and context, leading to better matching and more natural interactions with candidates.

What to look for: Systems that use modern LLMs, can understand context and nuance, and provide natural language interactions.

LIME (Local Interpretable Model-agnostic Explanations): A technical method for explaining individual AI decisions by showing which factors were most important for a specific recommendation. This helps understand the AI's reasoning for particular cases.

Real-world example: LIME might show that for a specific candidate, the AI's recommendation was based 40% on technical skills match, 25% on experience level, 20% on cultural fit indicators, 10% on location, and 5% on educational background.

Why it matters: LIME provides detailed explanations for individual decisions, helping recruiters understand why specific candidates were recommended or rejected.

What to look for: Systems that provide detailed explanations for individual recommendations, not just general explanations of how the system works.

M

Machine Learning (ML): A subset of AI that enables systems to learn and improve from experience without being explicitly programmed for every scenario. ML systems identify patterns in data and use these patterns to make predictions or decisions.

Real-world example: An ML system might analyse thousands of successful hires to learn that candidates who demonstrate certain communication patterns in their cover letters, have specific types of progression in their career history, and show particular combinations of skills tend to perform well in your organisation.

Why it matters: ML enables AI systems to improve over time and adapt to your organisation's specific needs and patterns.

What to look for: Systems that use appropriate ML techniques, can learn from your data, and improve performance over time.

Model Training: The process of teaching an AI system by showing it examples of good and poor recruitment decisions, allowing it to learn patterns and make future predictions. This is like teaching a new recruiter by showing them examples of successful and unsuccessful hires.

Real-world example: Training might involve showing the AI thousands of candidate profiles along with information about whether they were hired and how well they performed. The AI learns to identify patterns that predict success.

Why it matters: The quality of model training directly impacts AI performance. Poor training leads to poor recommendations.

What to look for: Comprehensive training data, robust training methodologies, and regular retraining to maintain accuracy.

Model Cards: Standardised documentation that describes how an AI system was built, its intended use, limitations, performance characteristics, and ethical considerations. These are increasingly required for procurement and compliance purposes.

Real-world example: A model card might detail that an AI recruitment system was trained on 500,000 candidate profiles from 2019-2024, achieves 85% accuracy on technical roles, has been tested for bias across gender and ethnicity, and is intended for initial screening only.

Why it matters: Model cards provide transparency about AI systems, help organisations make informed procurement decisions, and are increasingly required for EU AI Act compliance.

What to look for: Comprehensive model cards from vendors, clear documentation of system limitations, and evidence of ethical considerations in AI development.

Model Drift: When AI system performance degrades over time due to internal changes, software updates, or system modifications, rather than external market changes. This is distinct from concept drift, which is caused by external factors.

Real-world example: A system update might accidentally change how the AI weights certain factors, or database modifications might alter how candidate information is processed, leading to different (potentially worse) recommendations for similar candidates.

Why it matters: Model drift can cause sudden performance degradation that's harder to detect than concept drift, potentially leading to poor hiring decisions until the issue is identified and resolved.

What to look for: Systems with model drift detection, version control for all system components, and processes for quickly identifying and resolving performance issues.

Multi-modal AI: AI systems that can process and understand different types of information simultaneously, such as text (CVs), images (photos), audio (voice interviews), and structured data (assessment scores).

Real-world example: A multi-modal AI might analyse a candidate's written application, their performance in a video interview, their scores on technical assessments, and their social media presence to create a comprehensive evaluation.

Why it matters: Multi-modal AI can provide richer, more comprehensive candidate insights by considering multiple information sources.

What to look for: Systems that can handle various data types, integrate different information sources, and provide holistic candidate assessments.

N

Natural Language Processing (NLP): AI's ability to understand, interpret, and generate human language. This enables systems to read CVs, understand job descriptions, communicate with candidates, and extract meaning from text.

Real-world example: NLP might read a candidate's CV and understand that "led digital transformation initiatives" indicates leadership experience and technical skills, even though these terms aren't explicitly mentioned. It might also analyse the tone and content of candidate communications to assess cultural fit.

Why it matters: NLP enables AI systems to work with unstructured text data, which comprises most recruitment information (CVs, job descriptions, interview notes).

What to look for: Advanced NLP capabilities, ability to understand context and nuance, and support for multiple languages if needed.

Neural Network: A type of AI inspired by how the human brain processes information, using interconnected nodes (neurons) to identify patterns and make decisions. Neural networks are particularly good at handling complex, non-linear relationships in data.

Real-world example: A neural network might identify that candidates with a specific combination of educational background, early career choices, and skill development patterns tend to succeed in leadership roles, even if these relationships aren't obvious to human recruiters.

Why it matters: Neural networks can identify complex patterns that simpler AI methods might miss, potentially leading to better candidate insights.

What to look for: Systems that use appropriate neural network architectures for recruitment tasks and can explain their complex decision-making processes.

O

Optical Character Recognition (OCR): Technology that converts images of text (such as scanned CVs or certificates) into machine-readable text that AI systems can analyse. This is essential for processing documents that aren't in digital text format.

Real-world example: OCR might convert a scanned PDF CV into searchable text, allowing the AI to extract information about skills, experience, and qualifications for matching purposes.

Why it matters: OCR enables AI systems to process a wider range of candidate documents, including scanned or image-based submissions.

What to look for: High-accuracy OCR capabilities, support for various document formats, and ability to handle different languages and layouts.

Outcome Bias: The tendency to judge AI system quality based on individual results rather than the overall quality of the decision-making process. One poor outcome doesn't necessarily indicate system failure.

Real-world example: If an AI system recommends a candidate who performs poorly after hiring, this might be due to factors beyond the AI's control (changed role requirements, poor onboarding, personal circumstances) rather than poor AI performance.

Why it matters: Outcome bias can lead to poor evaluation of AI systems and knee-jerk reactions to individual cases rather than systematic assessment.

What to look for: Systems that track overall performance trends, not just individual outcomes, and provide context for understanding results.

P

Passive Candidate: Someone who isn't actively job hunting but might be open to the right opportunity. These candidates often represent the highest-quality talent pool but are harder to identify and engage.

Real-world example: AI might identify passive candidates by analysing their online activity, career progression patterns, and engagement with your company's content, then help craft personalised outreach messages that resonate with their interests and career goals.

Why it matters: Passive candidates often represent the best talent but require different engagement strategies than active job seekers.

What to look for: Systems that can identify passive candidates, provide insights into their motivations, and help craft effective engagement strategies.

Predictive Analytics: Using historical data, statistical algorithms, and machine learning to forecast future outcomes. In recruitment, this might predict which candidates are most likely to succeed, accept offers, or stay with the company long-term.

Real-world example: Predictive analytics might identify that candidates with certain educational backgrounds, specific types of early career experience, and particular skill combinations have an 85% likelihood of succeeding in senior management roles and a 92% probability of staying with the company for at least three years.

Why it matters: Predictive analytics helps make more informed hiring decisions, reduce turnover, and improve resource allocation.

What to look for: Proven predictive accuracy, relevant prediction models, and clear validation of predictive claims.

Precision: The percentage of AI recommendations that are actually correct. High precision means the system makes fewer mistakes when it identifies candidates as suitable, but it might miss some good candidates.

Real-world example: If an AI system recommends 100 candidates and 80 of them turn out to be genuinely suitable, the precision is 80%. Higher precision means fewer wasted interviews with unsuitable candidates.

Why it matters: Precision affects the efficiency of your recruitment process—low precision leads to time wasted on unsuitable candidates.

What to look for: Systems with high precision rates, ability to adjust precision/recall balance based on your needs, and clear measurement of accuracy.

Pre-screening: The initial filtering of candidates based on basic criteria before human review. AI can automate this process while ensuring consistency and fairness across all applications.

Real-world example: AI might automatically screen applications to identify candidates who meet minimum qualifications (education, experience, location), have required skills, and demonstrate basic cultural fit, then rank them for human review.

Why it matters: Automated pre-screening can significantly reduce recruiter workload while ensuring no candidates are overlooked due to human limitations.

What to look for: Flexible screening criteria, fair and consistent application, and clear audit trails for screening decisions.

Prompt Engineering: The practice of crafting effective instructions or queries for generative AI systems to produce desired outputs. In recruitment, this involves creating prompts that generate high-quality job descriptions, candidate communications, or interview questions.

Real-world example: Instead of asking AI to "write a job description for a marketing manager," effective prompt engineering might specify: "Write an inclusive job description for a senior marketing manager role in fintech, focusing on essential requirements, avoiding gendered language, and emphasising growth opportunities."

Why it matters: Good prompt engineering dramatically improves AI output quality, reduces hallucinations, and ensures generated content aligns with your organisation's needs and values.

What to look for: Systems with built-in prompt templates, guidance for effective prompting, and ability to customise prompts for your organisation's style and requirements.

Q

Quality Score: A numerical rating that AI systems assign to candidates based on how well they match job requirements and predicted success factors. These scores should be explainable, fair, and regularly validated.

Real-world example: A candidate might receive a quality score of 8.2/10 based on their skills match (9/10), experience level (8/10), cultural fit (7/10), and growth potential (8/10), with clear explanations for each component.

Why it matters: Quality scores help prioritise candidates and make objective comparisons, but they must be transparent and fair to be useful.

What to look for: Explainable scoring systems, regular validation against actual outcomes, and fair scoring across different demographic groups.

Query: A search request or question posed to an AI system. The sophistication of query handling affects how well the system can understand and respond to complex recruitment needs.

Real-world example: A query might be "Find software engineers with Python experience who have worked in fintech, are open to remote work, and have leadership potential"—the AI should understand all these requirements and their relationships.

Why it matters: Good query handling enables more precise candidate searches and better utilisation of AI capabilities.

What to look for: Systems that can handle complex, multi-faceted queries and provide relevant, ranked results.

R

Recall: The percentage of suitable candidates that the AI system successfully identifies. High recall means the system catches most good candidates, but it might also include some unsuitable ones.

Real-world example: If there are 50 genuinely suitable candidates among 1,000 applicants, and the AI identifies 40 of them, the recall is 80%. Higher recall means fewer good candidates are missed.

Why it matters: Recall affects your ability to identify top talent—low recall means you might miss excellent candidates.

What to look for: Systems with high recall rates, ability to balance recall with precision, and regular validation against actual hiring outcomes.

Recommendation Engine: AI systems that suggest candidates for roles, roles for candidates, or actions for recruiters to take, based on learned patterns and preferences. These systems aim to surface the most relevant options from large pools of possibilities.

Real-world example: A recommendation engine might suggest that based on a candidate's profile, they would be suitable for three specific open roles, recommend five candidates who would be perfect for a newly posted position, or suggest that a recruiter should prioritise contacting certain passive candidates this week.

Why it matters: Recommendation engines help manage information overload and ensure that good matches aren't missed due to the volume of data.

What to look for: Accurate recommendations, ability to explain reasoning, and systems that learn from feedback to improve over time.

Reinforcement Learning from Human Feedback (RLHF): A training method where AI systems learn to improve their recommendations by receiving feedback from human recruiters about the quality of their suggestions, gradually adapting to your organisation's preferences and success patterns.

Real-world example: When a recruiter marks an AI recommendation as "excellent match" or "poor fit," the system learns from this feedback to make better suggestions in future. Over time, it adapts to understand what your organisation values in candidates beyond what's captured in job descriptions.

Why it matters: RLHF enables AI systems to continuously improve and adapt to your organisation's unique culture and requirements, leading to increasingly accurate and relevant recommendations.

What to look for: Systems that can incorporate recruiter feedback, demonstrate improvement over time, and maintain learning capabilities without compromising fairness or introducing bias.

Regulatory Compliance: Ensuring AI systems meet all relevant legal requirements, including employment law, data protection regulations, and emerging AI-specific legislation. This is a complex and evolving area.

Real-world example: Your AI system must comply with UK employment law (fair recruitment practices), GDPR (data protection), the upcoming EU AI Act (high-risk AI systems), and any sector-specific regulations that apply to your industry.

Why it matters: Non-compliance can result in legal challenges, regulatory fines, and reputational damage. Compliance requirements are also evolving rapidly.

What to look for: Systems designed with compliance in mind, regular legal updates, and vendors who understand your regulatory environment.

Resume Parsing: The AI's ability to extract structured information from CVs and resumes, regardless of formatting, layout, or style. This involves understanding that the same information might be presented in many different ways.

Real-world example: AI might recognise that "Led team of 12 software developers" and "Team leader - 12 direct reports (software engineering)" both indicate leadership experience with team size, even though they're formatted completely differently.

Why it matters: Accurate resume parsing is fundamental to effective AI recruitment systems—if the AI can't properly extract information, it can't make good matches.

What to look for: High parsing accuracy across different formats, ability to handle creative or non-standard CVs, and extraction of both explicit and implicit information.

Return on Investment (ROI): The financial benefits of AI recruitment systems compared to their costs, typically measured through reduced time-to-hire, improved hire quality, decreased recruiter workload, and enhanced candidate experience.

Real-world example: An AI system costing £50,000 annually might deliver ROI through reducing time-to-hire by 15 days (saving £200,000 in opportunity costs), improving hire quality by 20% (reducing turnover costs by £150,000), and enabling recruiters to handle 30% more roles (avoiding two additional recruiter hires worth £120,000).

Why it matters: ROI measurement is essential for justifying AI investments, demonstrating value to leadership, and making informed decisions about system improvements and expansions.

What to look for: Clear ROI calculation methodologies, benchmarking against pre-AI performance, and systems that can track and report on key value drivers.

Right to Explanation: The legal requirement for individuals to understand how automated systems make decisions that affect them, including the right to receive clear explanations of AI recruitment decisions and the factors that influenced them.

Real-world example: If an AI system rejects a candidate's application, they have the right to understand why - for instance, that the decision was based on insufficient relevant experience (40% of decision), skills gap in required areas (35%), and location mismatch (25%).

Why it matters: Right to explanation is legally required under GDPR and other privacy laws, helps build trust with candidates, and enables organisations to demonstrate fair and transparent recruitment processes.

What to look for: Systems that provide clear, understandable explanations for all automated decisions, different explanation levels for different audiences, and compliance with relevant privacy regulations.

S

Scalability: How well an AI system can handle increased workload (more candidates, additional roles, extra locations) without significant performance degradation or cost increases.

Real-world example: A scalable system might handle 100 applications per month just as effectively as 10,000 applications per month and easily expand from one office to multiple global locations without requiring major reconfiguration.

Why it matters: Scalability affects your ability to grow and handle varying recruitment volumes. Poor scalability can lead to system failures during peak hiring periods.

What to look for: Proven scalability track record, cloud-based architecture, and transparent pricing models that don't penalise growth.

Sentiment Analysis: AI's ability to understand emotions, attitudes, and opinions expressed in text or speech. In recruitment, this might help assess candidate enthusiasm, cultural fit, or communication style.

Real-world example: Sentiment analysis might detect that a candidate's cover letter expresses genuine enthusiasm for the role and company mission or identify that their communication style aligns with your company culture.

Why it matters: Sentiment analysis can provide insights into candidate motivation and cultural fit that aren't captured by traditional CV screening.

What to look for: Accurate sentiment detection, cultural sensitivity, and appropriate use of sentiment information in decision-making.

Semantic Search: AI-powered search that understands the meaning and intent behind queries rather than just matching exact keywords. This enables more intuitive and effective candidate searches using natural language.

Real-world example: A semantic search might understand that searching for "team leadership experience" should return candidates who have "managed direct reports," "led project teams," or "supervised staff," even if they don't use the exact phrase "team leadership."

Why it matters: Semantic search dramatically improves the quality of candidate searches, reduces the time spent crafting complex Boolean queries, and helps identify qualified candidates who might be missed by keyword-only searches.

What to look for: Systems that support natural language queries, can understand synonyms and related concepts, and demonstrate improved search results compared to traditional keyword matching.

SHAP (SHapley Additive exPlanations): A method for explaining AI decisions by calculating how much each input factor contributed to the outcome. This provides detailed, quantitative explanations for individual recommendations.

Real-world example: SHAP might show that for a specific candidate recommendation, their technical skills contributed +0.3 to the score, relevant experience +0.2, leadership potential +0.1, but limited industry experience -0.1, resulting in a final recommendation score.

Why it matters: SHAP provides precise, quantitative explanations that help understand exactly why the AI made specific recommendations.

What to look for: Systems that provide detailed explanations for individual decisions, not just general explanations of system behaviour.

Skills Gap Analysis AI-powered identification of missing skills in candidate pools or current workforce, helping organisations understand talent market constraints and adjust recruitment strategies accordingly.

Real-world example: AI might analyse your candidate pipeline and identify that whilst there are plenty of junior Python developers available, there's a significant shortage of senior developers with machine learning expertise, prompting you to adjust salary expectations, expand geographic search, or develop internal training programmes.

Why it matters: Skills gap analysis helps set realistic hiring expectations, informs compensation strategies, and enables proactive talent planning for future needs.

What to look for: Systems that can analyse both internal talent and external market conditions, provide actionable insights about skill availability, and help develop strategies to address identified gaps.

Skills Taxonomy A structured database of skills and their relationships, helping AI systems understand connections between different competencies, technologies, and abilities. This enables more sophisticated matching beyond simple keyword searches.

Real-world example: A skills taxonomy might understand that "Python programming" is related to "data analysis," "machine learning," and "software development," and that someone with "Java" experience might quickly learn Python. It might also know that "project management" skills are valuable across industries.

Why it matters: Sophisticated skills taxonomies enable better candidate matching and help identify transferable skills that might be missed by simple keyword matching.

What to look for: Comprehensive, regularly updated skills databases, understanding of skill relationships, and ability to identify transferable competencies.

Shadow Banning/Silent Filtering: When candidates are effectively screened out, deprioritised, or disadvantaged without being formally rejected or notified. This can happen through algorithmic decisions that systematically reduce certain candidates' visibility or ranking.

Real-world example: An AI system might consistently rank candidates from certain postcodes lower in search results, or automatically filter out candidates with employment gaps, without explicitly rejecting them or explaining why they weren't considered.

Why it matters: Shadow banning raises serious fairness and transparency concerns, can constitute discrimination, and violates principles of candidate rights and procedural fairness.

What to look for: Systems with transparent ranking and filtering processes, clear candidate communication about screening decisions, and audit capabilities to detect silent filtering.

Supervised Learning: Training AI systems using examples where the correct answer is known, allowing the system to learn patterns and relationships. This is like teaching someone by showing them examples of correct and incorrect solutions.

Real-world example: Supervised learning might involve showing an AI system thousands of candidate profiles along with information about whether they were hired and how well they performed, allowing it to learn what characteristics predict success.

Why it matters: Supervised learning enables AI systems to learn from your organisation's specific experiences and patterns. *What to look for:* Systems that can learn from your historical data, appropriate training methodologies, and regular retraining to maintain accuracy.

T

Talent Pipeline: A pool of potential candidates for current and future roles, including both active applicants and passive candidates who have expressed interest. AI can help maintain and nurture these relationships automatically.

Real-world example: AI might maintain a talent pipeline by tracking candidate engagement with your company, identifying when passive candidates might be ready for new opportunities, and automatically nurturing relationships through personalised communications.

Why it matters: A strong talent pipeline reduces time-to-hire and ensures you have qualified candidates ready when roles become available.

What to look for: Systems that can maintain long-term candidate relationships, predict candidate readiness, and automate pipeline nurturing.

Time-to-Productivity: The period required for new hires to become fully effective in their roles, as predicted by AI systems based on candidate characteristics, role requirements, and historical performance data.

Real-world example: AI might predict that a candidate with extensive relevant experience will reach full productivity in 6 weeks, whilst a career-changer with strong transferable skills might require 12 weeks, helping you plan onboarding resources and set realistic expectations.

Why it matters: Time-to-productivity predictions help optimise hiring decisions, plan onboarding investments, and set appropriate expectations for new hire performance.

What to look for: Systems that can predict onboarding timelines, validate predictions against actual outcomes, and provide insights into factors that accelerate or delay productivity.

Training Data: The historical information used to teach AI systems, including past recruitment decisions, candidate profiles, hiring outcomes, and performance data. The quality and representativeness of training data directly impact AI performance.

Real-world example: Training data might include 10,000 candidate profiles with information about their backgrounds, the roles they applied for, whether they were hired, and how they performed if hired. The AI learns patterns from this data to make future predictions.

Why it matters: Training data quality is crucial for AI effectiveness. Biased or poor-quality training data leads to biased or inaccurate AI recommendations.

What to look for: Large, diverse, high-quality training datasets, regular data updates, and processes for identifying and correcting data quality issues.

Transfer Learning: Using knowledge gained from one domain or task to improve performance in another related area. This is particularly useful for organisations with limited historical data.

Real-world example: An AI system might use general knowledge about successful project managers to help evaluate candidates for a specific project management role, even if there's limited historical data about that position.

Why it matters: Transfer learning enables AI systems to work effectively even with limited domain-specific data, making them more practical for diverse organisations.

What to look for: Systems that can leverage broader knowledge, pre-trained models, and ability to adapt general knowledge to specific contexts.

Transformer: A type of AI architecture that's particularly effective at understanding language and context. Transformers power many modern language models and are crucial for sophisticated text analysis in recruitment.

Real-world example: A transformer-based system might understand that "led digital transformation initiatives" indicates both leadership skills and technical aptitude, and that this experience is relevant for various senior roles across different industries.

Why it matters: Transformer-based systems can provide more sophisticated language understanding, leading to better CV analysis and candidate matching.

What to look for: Systems that use modern transformer architectures, can understand context and nuance, and provide sophisticated language analysis.

U

Unconscious Bias: Hidden preferences or prejudices that affect decision-making without our conscious awareness. These biases can influence recruitment decisions and may be reflected in AI systems trained on biased data.

Real-world example: A recruiter might unconsciously favour candidates from prestigious universities or with traditionally "Western" names, even when these factors aren't relevant to job performance. AI systems can either help reduce these biases or amplify them if not carefully designed.

Why it matters: Unconscious bias can lead to discrimination and missed opportunities. AI systems can help identify and reduce bias, but they can also perpetuate it if not properly managed.

What to look for: Systems designed to identify and mitigate bias, regular bias testing, and processes for recognising and addressing unconscious preferences.

Unsupervised Learning: AI learning that identifies patterns in data without being told what to look for. This can discover hidden insights and relationships that humans might miss.

Real-world example: Unsupervised learning might identify that candidates with certain combinations of experiences, even if they seem unrelated, tend to succeed in specific roles. It might discover that successful sales people often have unexpected backgrounds like music or sports.

Why it matters: Unsupervised learning can reveal valuable insights about candidate success factors that aren't obvious to human recruiters.

What to look for: Systems that can discover meaningful patterns, validate discovered insights, and translate findings into actionable recommendations.

User Experience (UX): How easy, efficient, and pleasant it is for recruiters, hiring managers, and candidates to interact with the AI system. Good UX is crucial for successful adoption and effective use.

Real-world example: Good UX might include intuitive dashboards that show key information at a glance, simple processes for reviewing candidate recommendations, clear explanations of AI decisions, and mobile-friendly interfaces for on-the-go access.

Why it matters: Poor UX leads to low adoption rates, user frustration, and reduced effectiveness. Even powerful AI systems are useless if people don't want to use them.

What to look for: Intuitive interfaces, user-centric design, comprehensive user testing, and ongoing UX improvements based on feedback.

V

Validation: The process of testing AI systems to ensure they work correctly, fairly, and effectively before full deployment. This includes testing accuracy, fairness, and performance under various conditions.

Real-world example: Validation might involve testing the AI system with historical data to see if it would have made good recommendations, testing for bias across different demographic groups, and verifying that explanations match actual decision-making processes.

Why it matters: Proper validation helps identify problems before they affect real recruitment decisions, ensuring that AI systems are effective and fair.

What to look for: Comprehensive validation processes, third-party validation, and ongoing validation as systems evolve.

Vector Databases: Specialised storage systems that enable fast similarity searches for candidate matching and skills comparison by storing AI-generated numerical representations of candidates, jobs, and skills.

Real-world example: A vector database might store numerical representations of all candidates' skills and experiences, enabling instant searches for candidates similar to your best performers or rapid identification of candidates with complementary skills to existing team members.

Why it matters: Vector databases enable more sophisticated matching capabilities, faster search performance, and more nuanced understanding of candidate-job fit beyond simple keyword matching.

What to look for: Systems that use modern vector storage technologies, demonstrate fast search performance, and can handle large volumes of candidate data efficiently.

Version Control: The practice of tracking different versions of AI models and being able to revert to previous versions if problems arise. This is crucial for maintaining system stability and managing updates.

Real-world example: If a new version of the AI system performs poorly or introduces bias, version control allows you to quickly revert to the previous, working version while the issues are resolved.

Why it matters: Version control provides safety nets for AI updates and helps manage the evolution of AI systems over time.

What to look for: Robust version control systems, easy rollback capabilities, and clear processes for managing updates.

Video Interview Analysis: AI that can analyse recorded video interviews for various factors such as communication skills, enthusiasm, cultural fit, or technical competence. This requires careful consideration of bias and privacy implications.

Real-world example: AI might analyse a video interview to assess communication clarity, enthusiasm level, and response quality, while being careful not to discriminate based on accent, appearance, or cultural communication styles.

Why it matters: Video analysis can provide insights into candidate qualities that aren't captured in CVs, but it also raises significant bias and privacy concerns.

What to look for: Systems that focus on job-relevant factors, extensive bias testing, privacy protections, and transparency about what's being analysed.

W

Workflow Automation: Using AI to handle routine recruitment tasks automatically, such as sending acknowledgement emails, scheduling interviews, or updating candidate status. This frees up recruiters for more strategic work.

Real-world example: Workflow automation might automatically send personalised acknowledgement emails to applicants, schedule interviews based on availability, send reminder emails before interviews, and update candidate records with interview feedback.

Why it matters: Workflow automation reduces manual work, ensures consistency, and improves candidate experience through faster responses.

What to look for: Flexible automation rules, integration with existing systems, and ability to personalise automated communications.

Workforce Planning: Using AI to predict future hiring needs based on business growth, employee turnover, market trends, and strategic objectives. This enables proactive recruitment strategies.

Real-world example: AI might predict that based on your growth plans and historical turnover patterns, you'll need to hire 15 software engineers, 8 sales representatives, and 3 marketing managers over the next 12 months, with specific timing recommendations.

Why it matters: Workforce planning helps ensure you have the right people at the right time, avoiding both understaffing and overstaffing.

What to look for: Accurate prediction models, integration with business planning, and ability to scenario-plan for different growth trajectories.

X

eXplainable AI (XAI): AI systems that can provide clear, understandable explanations for their decisions and recommendations. This is crucial for building trust, ensuring fairness, and meeting regulatory requirements.

Real-world example: An explainable AI might say "I recommended Sarah for this role because she has 8 years of relevant experience (weighted 30%), strong technical skills in required areas (25%), demonstrated leadership (20%), and good cultural fit indicators (15%). Her limited experience in our specific industry reduced the score by 5%."

Why it matters: Explainability is essential for legal compliance, building user trust, and enabling informed decision-making.

What to look for: Clear, specific explanations, different explanation levels for different users, and validation that explanations accurately represent system reasoning.

Y

Yield Rate: The percentage of candidates who accept job offers. AI can help predict and improve yield rates by better matching candidates to roles and company culture.

Real-world example: AI might identify that candidates with certain motivations and career goals are more likely to accept offers, helping you prioritise outreach and tailor offer presentations to improve acceptance rates.

Why it matters: Higher yield rates reduce recruitment costs and time-to-hire by ensuring fewer declined offers.

What to look for: Systems that can predict offer acceptance likelihood, provide insights into candidate motivations, and help optimise offer strategies.

Z

Zero-shot Learning: An AI's ability to handle new situations without specific training examples. This enables systems to work with novel job types or unusual candidate profiles.

Real-world example: A zero-shot learning system might be able to evaluate candidates for a completely new role type by understanding the general requirements and applying knowledge from similar roles.
Why it matters: Zero-shot learning makes AI systems more flexible and able to handle the variety of real-world recruitment scenarios.

What to look for: Systems that can generalise to new situations, handle unusual cases gracefully, and adapt to novel requirements.

Quick Reference: Key Metrics Comparison

Metric	Focus	Best For	Limitation
Precision	Accuracy of positive predictions	Reducing wasted interviews	May miss good candidates
Recall	Catching all suitable candidates	Ensuring no talent is missed	May include unsuitable candidates
Demographic Parity	Equal selection rates across groups	Basic fairness checking	May conflict with accuracy
Equalised Odds	Equal accuracy across groups	Sophisticated fairness	Complex to implement
Individual Fairness	Similar treatment for similar candidates	Personalised fairness	Difficult to define similarity

Mini Vendor Evaluation Checklist

Ethics & Fairness:

- [ ] Regular third-party bias audits
- [ ] Multiple fairness metrics monitoring
- [ ] Clear bias mitigation strategies
- [ ] Transparent decision explanations

Performance & Reliability:

- [ ] Validated accuracy claims
- [ ] Concept and model drift detection
- [ ] Performance monitoring dashboards
- [ ] Rollback capabilities for failed updates

Compliance & Governance:

- [ ] GDPR compliance framework
- [ ] EU AI Act preparation
- [ ] Comprehensive audit trails
- [ ] Model cards/documentation

Technical Capabilities:

- [ ] Modern AI architecture (transformers, LLMs)
- [ ] Integration with existing systems
- [ ] Scalability evidence
- [ ] Hallucination detection (for generative AI)

Use Case Index

If you're using AI for sourcing, read:

- Boolean Search, Passive Candidate, Talent Pipeline, Skills Taxonomy

If you're using AI for screening, read:

- Resume Parsing, Candidate Matching, Pre-screening, Quality Score

If you're using generative AI, read:

- Large Language Model, Generative AI, Hallucination, Prompt Engineering

If you're concerned about bias, read:

- Algorithmic Bias, Fairness Metrics, Intersectionality, Shadow Banning

If you're preparing for compliance, read:

- GDPR, EU AI Act, Audit Trail, Model Cards

If you're concerned about data security, read:

- Data Poisoning, GDPR, External Data Sources

If you're presenting to the board, read:

- AI Governance, Predictive Analytics, Key Performance Indicators

If you're implementing governance frameworks, read:

- AI Governance, Model Cards, Human Oversight, Regulatory Compliance

Regulatory Landscape: EU AI Act Impact

High-Risk Classification

Recruitment AI is classified as **high-risk** under the EU AI Act, requiring:

- Conformity assessments before deployment
- Detailed documentation and record-keeping
- Human oversight for significant decisions
- Transparency obligations to candidates
- Risk management systems
- Accuracy and robustness testing

Implementation Timeline

- **2024**: Prohibited AI practices banned
- **2025**: High-risk AI system requirements (including recruitment)
- **2026**: General-purpose AI model requirements
- **2027**: Full implementation and enforcement

UK Implications

While not directly applicable, UK companies should consider EU AI Act compliance if they:

- Process EU candidate data
- Have EU subsidiaries
- Serve EU clients
- Want to maintain EU market access

Key EU AI Act Requirements for Recruitment AI:

- **Risk Assessment**: Comprehensive evaluation of potential harms
- **Data Governance**: Quality standards for training and testing data
- **Documentation**: Detailed technical documentation and model cards
- **Human Oversight**: Meaningful human control over AI decisions
- **Transparency**: Clear information to candidates about AI use
- **Accuracy & Robustness**: Testing and validation requirements
- **Bias Monitoring**: Regular assessment for discriminatory impacts

Privacy by Design Principles

Modern AI recruitment systems should incorporate privacy protection from the ground up:

Data Minimisation: Collect only necessary candidate information **Purpose Limitation**: Use data only for specified recruitment purposes **Storage Limitation**: Retain data only as long as needed **Accuracy**: Ensure candidate data is accurate and up-to-date **Security**: Protect data with appropriate technical measures **Transparency**: Clearly communicate data use to candidates **Accountability**: Demonstrate compliance with privacy principles

Risk Management Framework

High-Risk Areas to Monitor:

- Algorithmic bias and discrimination
- Privacy breaches and data security
- Regulatory compliance failures
- Candidates experience degradation
- System performance issues
- Vendor dependency risks

Mitigation Strategies:

- Regular bias audits and testing
- Comprehensive security measures
- Legal compliance monitoring
- Continuous performance assessment
- Vendor risk management
- Incident response procedures

Terms for Senior Leadership

When presenting to senior leadership, these terms are particularly important:

Algorithmic Bias: The risk that AI systems might unfairly discriminate against certain groups, potentially leading to legal challenges and missed talent opportunities.

Human-in-the-Loop: Ensuring human judgement remains central to important decisions, with AI providing support rather than replacement.

GDPR Compliance: Meeting European data protection requirements, with significant financial penalties for non-compliance.

Predictive Analytics: Using AI to forecast recruitment outcomes, helping make more informed decisions about candidates and hiring strategies.

Explainable AI: The ability to understand and justify AI decisions, crucial for legal compliance and stakeholder confidence.

Scalability: How well the AI system can grow with your organisation without performance degradation or prohibitive costs.

Common Misconceptions Explained

"AI will replace recruiters" Reality: AI augments human capabilities rather than replacing them. The most effective systems enhance human judgement, automate routine tasks, and provide insights that help recruiters make better decisions. Human skills like relationship building, cultural assessment, and strategic thinking remain irreplaceable.

"AI is completely objective" Reality: AI systems can perpetuate or amplify existing biases if not carefully designed and monitored. They learn from historical data that may contain human biases, and their algorithms can introduce new forms of discrimination. Regular bias testing and mitigation strategies are essential.

"AI is too complex for HR professionals to understand" Reality: While the technical details are complex, modern AI systems are designed with user-friendly interfaces. You don't need to understand the underlying algorithms to use them effectively, just as you don't need to understand combustion engines to drive a car.

"AI always makes better decisions than humans" Reality: AI excels at processing large amounts of data quickly and identifying patterns, but human judgement remains crucial for complex decisions, cultural fit assessment, and handling unusual situations. The best outcomes come from AI-human collaboration.

"AI systems are 'set and forget'" Reality: AI systems require ongoing monitoring, regular updates, and continuous improvement. They need feedback to learn, retraining to stay current, and oversight to ensure they remain fair and effective.

"More data always leads to better AI" Reality: Data quality is more important than quantity. Biased, incomplete, or outdated data can lead to poor AI performance regardless of volume. Clean, representative, current data is essential for effective AI systems.

Industry-Specific Considerations

Financial Services: Particular attention to regulatory compliance, risk management, and security clearance requirements.

Healthcare: Focus on professional licensing verification, patient safety considerations, and complex qualification requirements.

Technology: Emphasis on rapidly evolving skill requirements, technical assessment integration, and global talent mobility.

Manufacturing: Consideration of safety requirements, practical skills assessment, and shift work compatibility.

Retail: Focus on customer service skills, seasonal hiring patterns, and location-specific requirements.

Education: Attention to teaching qualifications, safeguarding requirements, and value-based cultural fit.

This glossary represents the current state of AI recruitment technology and terminology. As the field continues to evolve rapidly, regular updates ensure it remains relevant and useful for HR professionals navigating the AI landscape.