

L3A WORKING GROUP

Name the Job Before You Trust It: Understanding AI Layers and Its Applications in HR

October 2026

Contributors

This report was produced by the following members of the L3A AI Working Group:

Working Group Members

- ☐ **Nicolas Douet – Deloitte**
- ☐ **Isabelle Hanin Dricot - IQEQ**
- ☐ **Giuliana Arcangeli - PWC**
- ☐ **Ingo Gozemba – Hauck Aufhaeuser**

Coordinators

- ☐ **Marta Bawor – L3A**

About L3A

The Luxembourg Alternative Administrators Association (“L3A”) is the leading representative body for fund and corporate administrators of alternative assets in Luxembourg.

Since its founding in 2004, the association has been dedicated to promoting the interests of alternative service providers and contributing to the development of Luxembourg as the investment hub and operations centre of choice for global alternative investment managers.

Our members deliver our mission through their work on our 5 permanent committees, being Operations, Tax, Regulatory, HR and Events & Communications, with output delivered to industry players, regulators and government bodies amongst others, and published through our website and social media channels, in particular LinkedIn.

As the recognised voice of our profession, we engage with authorities through participation in commissions and working groups, and maintain strong ties with government bodies, professional organisations, and Chambers of Commerce.

Our membership is made up of over 60 members, including Full Members, being Luxembourg administration firms predominantly providing services to foreign investment managers of funds and related entities in the alternative management industry and Associate Members, being Luxembourg professional firms operating in the real asset environment and offering services to Full Members and/or their clients.

Together, we form a trusted network committed to excellence.

Introduction

A practical lens (diagnostic, predictive, generative) that tells HR teams what a tool is worth, what it risks, and how to use it well. With a Luxembourg reality check.

AI is a spectrum of capabilities that have evolved from basic analytics to advanced generative systems. AI enables machines to perform tasks that traditionally require human intelligence, such as reasoning, learning, problem-solving, and language understanding.

Luxembourg is not an AI laggard. In 2025, roughly one in three Luxembourg enterprises used at least one AI technology, the fifth-highest share in the European Union and comfortably ahead of the EU average of one in five (Eurostat).

For HR, that figure has a sharp practical edge: the AI is already in your building. It is screening or ranking applicants in the recruitment system, suggesting courses in the learning platform, drafting messages in the collaboration suite, and surfacing trends on the analytics dashboard.


The pressure to "do something with AI" tends to arrive at the same moment as a quieter, more awkward problem: the word "AI" has been stretched to mean almost anything, from a rules-based chatbot following a fixed decision tree to a generative model writing a paragraph from scratch.

That vagueness is not harmless. It produces two expensive mistakes in opposite directions: buying sophisticated, higher-risk AI to do a job a simple filter would have done better and more safely, or leaving genuine capability switched off because no one realized it was there. In both cases the root problem is the same: the team could not say, in plain language, what the tool was actually doing.

You do not need a data-science background to fix this. You need one habit, and it fits on a postcard.

Three jobs, Three Questions

Almost every application of AI in HR is doing one of three jobs, and each job answers a different question.



The same problem can call on all three: when turnover jumps in a team, a diagnostic tool helps you understand *why* it happened, a predictive tool flags *who* might be next, and a generative tool helps you *act* by drafting a manager check-in guide. Same situation, three different jobs, three different questions answered. None of these is more "advanced" than the others in any way that matters to you. What matters is that the job a tool is doing determines three things at once: how much value it can realistically deliver, how much risk it carries, and what controls you need to wrap around it.

The same lens, applied to a Luxembourg hiring desk

Picture a mid-sized Luxembourg employer, a fund administrator, say, filling a compliance role. The advertisement runs in English; applications arrive in English, French, and German, from residents and from the cross-border workers who make up nearly half the country's workforce. (In 2024, 47% of Luxembourg's 489,000 employees commuted in, most from France, and only one in four employees was a Luxembourg national.) The recruitment system offers a feature that scores and ranks each CV against the role. It is tempting to switch it on and trust the order it produces.

Run the three questions first.

Is it generative? In part, if it also drafts rejection emails or translates the posting, that is production, and it needs human review before anything reaches a candidate.

Alright, but is it diagnostic? Not in the ranking itself: a score attached to a CV explains nothing about the past, it forecasts fit. But there is a diagnostic job hiding in the same data, and it is the safest place to start. Turned backward, the tool can answer real questions: why this role has been hard to fill, where applicants abandon the process, what the firm's past compliance hires actually have in common. That work touches only historical, aggregated data and affects no live candidate, so the risk is modest and the controls can stay light. It also puts the firm's own hiring patterns in plain view, and those patterns are precisely what a predictive tool would learn from.

So, is it predictive? Yes, at its core: a ranking is a forecast of fit, and that is exactly where the risk concentrates. Here the Luxembourg context bites. A model trained on whom the firm has hired before can quietly learn that a particular university, a particular school system, or fluency in one language "predicts" success, and start down-ranking equally capable candidates who trained in Metz, Trier, or Arlon rather than in a familiar institution. None of those attributes is a legitimate selection criterion. Several are proxies for nationality and Luxembourg's Labour Code (Article L.251-1) explicitly prohibits discrimination, direct or indirect, on the basis of real or supposed nationality, with the Inspection du Travail et des Mines empowered to enforce it.

That single reframing: *"this is a predictive tool, and prediction is where bias hides"* changes what you do next. It tells you this is not a feature to switch on and trust, but a system that needs bias testing, a human firmly in the loop, and two specific Luxembourg steps before it ever runs: prior consultation of the staff delegation, required under the Labour Code where one exists, and a clear notice to candidates that an automated tool is involved (a transparency duty under both the GDPR and the EU AI Act). Article 22 of the GDPR adds a hard limit: no candidate may be rejected by the algorithm alone, a person must decide.

What good looks like for the organization

For an HR leadership team, the three-question lens turns several hard problems into manageable ones. It makes an honest inventory possible. Walk through the AI features already present in your systems and label each one: is it explaining the past, predicting the future, or producing content? This is usually a revealing exercise. It surfaces predictive scoring that has been running quietly with no human review, and it surfaces generative or analytical capability that has been paid for and never switched on. You cannot govern or exploit what you have not named.

It sharpens how you buy. Make one question standard in every vendor conversation: *which of the three jobs is this system doing, and is that level of sophistication necessary for the value we want?* That single question exposes "AI-washing" (a deterministic rules engine dressed up as artificial intelligence) and it exposes over-engineering, where a generative model is being sold for a task, a structured template would handle more predictably and more cheaply.

And it lets you match oversight to the job rather than applying one blunt policy to everything. Generative drafting needs human review before anything reaches a person, plus some discipline about how prompts are written. Predictive systems that score or rank people need bias testing and a human in the decision loop. Diagnostic analysis of aggregated, anonymized data typically needs the lightest touch. The lens does not replace a governance framework, but it tells you where to point one.

Employee Trust, Communication & Change Management

None of this lands without trust

A tool that is technically sound and legally compliant can still fail if the people it affects do not trust the process. Indeed, 37 % of the global workforce trusts businesses to deploy AI in a way that is good for them, not for shareholders, not for efficiency, but for them personally (Edelman, 2025). In Luxembourg, where social dialogue is not optional but structural, staff delegations have legally codified information, consultation, and in some cases co-decision rights under the Labour Code, ignoring the trust question is not just unwise, it is procedurally risky.

Three principles help:

Explain before you deploy. Tell employees and candidates, in plain language, which AI features are active, what job each one is doing (diagnostic, predictive, or generative), and what role a human plays in every decision. The OECD's AI Principles list transparency and explainability as foundational requirements for trustworthy AI, and the EU AI Act's own transparency obligations (Articles 13 and 50) operationalise the same idea (OECD, 2024 revision; EU AI Act). In practice, this means a short internal communication, not a privacy notice buried in an annex, before any new AI feature goes live.

Involve the staff delegation early, not late. Luxembourg's Labour Code requires prior consultation where automated monitoring or decision-support tools are introduced. But the minimum legal requirement is a floor, not a ceiling. Involving employee representatives in the inventory exercise described above, letting them see the labels, ask questions, and flag concerns, converts a compliance step into a credibility signal. Research by the CIPD confirms that employee involvement in AI adoption decisions is one of the strongest predictors of perceived fairness and sustained adoption (CIPD, *People Profession 2030*, 2024).

Treat change management as a workstream, not an afterthought. AI in HR is not a single technology deployment; it is a rolling series of capability changes that alter how recruiters screen, how managers review, and how employees experience the organisation. Each change needs its own answer to three questions employees will always ask: *What is changing for me? Why? And what support do I get?* Organisations that answer those questions deliberately, through structured change management rather than a launch email, see materially higher adoption and lower resistance (Prosci, *Best Practices in Change Management*, 12th ed., 2023).

The three-question lens helps here too. Telling a recruiter "we have added AI to the recruitment system" is vague and slightly threatening. Telling them "the system now has a diagnostic feature that shows where applicants drop out, and a generative feature that drafts first-pass job ads for your review" is concrete, bounded, and honest about the human role. Name the job, and the change becomes easier to explain and easier to trust.

What good looks like on a Monday morning

Use AI with judgment

Five rules to remember before you trust an AI output

Before you act on any result, identify what kind of AI work is happening — and how much confidence it deserves.

- 1. Name the job first**
 Before you use an output, label it — even silently: generating a draft, predicting a risk, or explaining a trend. That one second changes how much weight you give the result.
 [Generate draft] [Predict risk] [Explain trend]
- 2. Generative AI = first draft**
 NEVER A DECISION
 A model can write the job ad, summary, or translation. You still own the words that go to a real person. Edit and approve — do not just forward.
- 3. Predictive AI = hypothesis**
 NEVER A VERDICT
 A risk score or ranking is a signal to look closer and have a conversation. The number starts the human work; it does not end it.
- 4. Diagnostic AI = question**
 NEVER THE ANSWER
 A pattern such as falling engagement tells you where to investigate. It does not tell you the cause. The tool points; you investigate.
- 5. Ask if AI is necessary at all**
 Sometimes a filter, checklist, or template will do the job with less risk and more consistency. Choosing the simpler tool is judgment, not a lack of ambition.

Label it. Review it. Investigate it. Decide like a human.
 THE OUTPUT IS SUPPORT — NOT AUTHORITY.

1. **Name the job before you use the output.** Say it, even silently: "this is generating a draft," "this is predicting a risk score," "this is explaining a trend." That one second of labeling changes how much weight you give the result.

2. **Treat generative output as a first draft, never a decision.** A model can write the job ad or the summary; you own the words that go to a real person. Edit and approve **but do not forward**.

3. **Treat predictive output as a hypothesis, never a verdict.** A flight-risk score or a ranking is a reason to look closer and have a conversation **but not a reason to act on someone**. The number is the beginning of the human work, not the end of it.

4. **Treat diagnostic output as a question, never an answer.** "Engagement dropped in this team" tells you where to look; **it does not tell you why**. The tool points; you investigate.

5. **Ask whether this level of AI is even necessary.** Sometimes a filter, a checklist, or a template would do the job with less risk and more consistency.

Choosing the simpler tool is a sign of judgment, not a lack of ambition.

These habits cost nothing and require no technical training.

They are the difference between a team that uses AI and a team that is used by it.

In Luxembourg, the regulator is closer than you think

Two features of the local landscape make the three-question habit more than good hygiene.

First, supervision is consolidating in a familiar place. Under Luxembourg's draft implementing law for the AI Act, the CNPD is set to act as the coordinating market surveillance authority and single point of contact across sectoral regulators, with the CSSF leading for supervised financial institutions (and now maintaining its own dedicated AI guidance).

The body that reviews your data processing will be the body that looks at your HR AI. And it has signalled where it is looking: in its guidance on AI literacy, the CNPD specifically noted that some employers already train HR staff to recognize the biases that can creep into a partly automated recruitment process. That is a regulator telling you, in advance, what good practice looks like.

That observation points to something the AI Act has now made explicit. Article 4 of the Regulation imposes an **AI literacy obligation** on every organisation that deploys or operates an AI system: providers and deployers must ensure that their staff have a "sufficient level of AI literacy," taking into account the technical knowledge, experience, education, and context of the persons involved (EU AI Act, Article 4). This is not a suggestion; it is a directly applicable requirement, in force since 2 February 2025.

For HR, "sufficient" has a practical meaning tied directly to the three-question lens. A recruiter using a predictive ranking tool needs to understand what a prediction is, why it can embed bias, and what their role is in the decision loop, not how the algorithm's weights are calibrated. A learning specialist using a generative tool to draft course descriptions needs to understand that the output is a first draft requiring review, not a vetted source of truth. Training should match the job the tool is doing, not deliver a generic "AI awareness" module that leaves people impressed but no more capable.

The CNPD's own AI literacy guidance recommends that training be **role-specific, documented, and refreshed** as tools evolve (CNPD, 2025). Concretely, this means:

- **For recruiters and HR advisors** using predictive or generative features: scenario-based training on bias recognition, output verification, and the obligation to keep a human in the decision loop.
- **For managers** receiving AI-generated insights (flight-risk scores, engagement dashboards, suggested development actions): training on what the numbers represent, what they do not, and why a score is the start of a conversation, never the end of one.
- **For HR leaders and procurement teams:** training on the three-question framework itself — enough to interrogate vendors, label features accurately, and match oversight to the job.

Document what you deliver. Under the AI Act's high-risk provisions (Articles 9 and 29), deployers of high-risk systems, and AI used in employment decisions qualifies, must be able to demonstrate that the humans exercising oversight are competent to do so. A training register, even a simple one, converts good practice into evidence.

Second, one category is not a risk to be managed but a line not to cross. Tools that claim to read employees' moods or emotions can look, at a glance, like sophisticated diagnostics. Under Article 5 of the

EU AI Act, AI systems that infer the emotions of people in the workplace are prohibited outright, with narrow exceptions only for medical or safety reasons, a ban already in force.

"Engagement analytics" built on aggregated, anonymized survey data is one thing; a system inferring an individual's emotional state is another, and it is off the table. Affect detection dressed up as "diagnostic insight" is the one place where the honest answer is not "govern it carefully" but "do not use it."

This is also why the most dangerous tools are the ones that quietly change jobs. An "interview summarizer" is generative and assistive, until its summaries start shaping who advances, at which point it is functioning as a predictive, determinative system that was never labeled as one. Judge a tool by what it actually influences, not by what the box says. The label is marketing; the job is what the tool does to real decisions about real people.

What to do this week

The three questions are only useful as a habit, so make them one, at both levels.

For HR leaders: run a one-page inventory of every AI feature already live in your stack, and label each one diagnostic, predictive, or generative. The exercise usually surfaces predictive scoring running without oversight, and paid-for capability sitting unused. Wherever you find a predictive tool touching hiring, promotion, or performance, treat it as high-risk under the EU AI Act: confirm the staff delegation has been consulted, that candidates and employees are told an automated tool is involved, and that a human owns every consequential decision. Then plug into the local conversation rather than navigating it alone — the CNPD and the Luxembourg AI Factory jointly run RE.M.I., a community of practice that gives HR a direct line to how regulators and peers are reading these questions in real time.

For everyone using the tools: adopt the one-second reflex. Name the job before you trust the output. "This is generating a draft", so I edit and own it. "This is predicting a score", so I treat it as a hypothesis and have the conversation. "This is explaining a trend", so I go and investigate.

Diagnostic, predictive, generative; why did this happen, what might happen next, can you produce this for me. The questions require no expertise, only the discipline to ask them, every time, before the output earns your trust.

Sources and further reading: Eurostat, "Use of artificial intelligence in enterprises" (2025 reference year); STATEC labour-market figures (2024); the EU AI Act (Regulation (EU) 2024/1689), notably the Article 5 prohibition on workplace emotion-recognition systems and the classification of employment-related AI as high-risk under Annex III; the Luxembourg Labour Code (Code du travail), Article L.251-1 on non-discrimination and the staff-delegation information and consultation provisions; and published guidance from the CNPD and the CSSF. Specific obligations and timelines remain subject to ongoing EU and national implementation; verify Labour Code article references and the current position with counsel before publication.