Human Level

Human Rights Due Diligence in the Age of AI

A Discussion Paper on the Potential Uses and Open Questions for AI in HRDD

September 2026

Contents

Introduction

Who this paper is for – and what you will get from it

This discussion paper is for practitioners, experts, civil society, trade unions, worker organisations, consultants, service providers, academics and others working on, implementing or advising on human rights due diligence (HRDD) / corporate sustainability due diligence under the EU Corporate Sustainability Due Diligence Directive (CSDDD).

It provides a practical structure for thinking about the growing use of artificial intelligence (AI) in HRDD: where AI could be used, what opportunities this might create, and what questions we should be asking as its use expands.

It is intended as a starting point for further discussion, rather than a set of answers.

Why this conversation matters now

AI is here. Its capabilities are expanding rapidly, and companies are increasingly exploring how it can support different components of HRDD.

We hear about AI being used to analyse large volumes of information, identify patterns and potential risks, support supply chain mapping, review grievances and audit findings, monitor developments, draft disclosures and help teams manage increasingly complex due diligence processes.

This use of AI in HRDD is likely to grow. Expectations of companies to conduct effective HRDD are becoming more firmly embedded in regulation and stakeholder expectations. And at the same time, AI systems are becoming more capable, more accessible and increasingly integrated into the tools and systems companies already use.

This brings us to a central question:

Can AI help strengthen the quality, scale and effectiveness of HRDD? Or does it risk accelerating superficial, compliance-driven and overly technocratic approaches that weaken meaningful engagement, accountability and outcomes for people?

We are at an important juncture.

Decisions being made now – about where AI is introduced, what it is asked to do, what information it relies on, where humans remain involved, and how its outputs are questioned and tested – will shape how HRDD evolves in practice.

Objective of this paper

Over the past year, through conversations with company practitioners, experts and advisors, we have heard both a growing range of ways in which AI is being used for HRDD, as well as a growing number of questions about what that use means in practice.

This paper brings those conversations together and provides a practical framework for interrogating the use of AI against the expectations of HRDD itself.

It works through the different elements of the due diligence process and asks: Where could AI potentially help? What could it enable? Where could it create pitfalls? What questions could practitioners be asking?

We have found that grounding the discussion on AI in the actual expectations of HRDD, a number of which are in turn reflected in the EU CSDDD, helps move the conversation from the abstract to the practical: whether, where and how it can contribute to more effective due diligence and better outcomes for people.

Our intention is not to advocate for or against the use of AI in HRDD, nor to provide definitive answers. It is to provide a structure for a shared and evolving conversation – one that can help companies make more informed choices about how they use AI, ask better questions of AI providers, and identify the safeguards and human judgment needed along the way.

This discussion paper delves into possible AI uses (potential opportunities), as well as questions to consider (food for thought) as we use AI in this way. We go through this use of AI for each component of HRDD, as well as policy and embedding. This discussion paper then concludes with eight themes that have surfaced consistently across our discussions on AI and HRDD regardless of which stage of the due diligence process we are considering.

To be clear, by referencing a possible use of AI in this paper, we are not advocating for this use or suggesting this is a helpful use of AI for HRDD. Rather, we are surfacing that this is a use that practitioners are already seeing as an opportunity, or are likely to see as an opportunity moving forward, which can in turn inform a discussion related to how to ensure this use contributes to meaningful HRDD in practice.

How you can contribute

This is a discussion paper, and we want the discussion to shape what comes next.

We will continue to build on this paper based on the experiences, perspectives and questions we hear from others working on AI and HRDD.

What we hear will feed into the next iteration of this work – helping us collectively build a clearer picture of where AI can strengthen HRDD, where the pitfalls lie, and what good practice could look like.

There are two ways you can contribute:

1. Share written feedback

We welcome reflections on any part of this paper, as well as additional examples, experiences, questions or perspectives that you think should be part of the discussion.

We can keep your contribution confidential if you would prefer, but in the interests of an open and transparent discussion, we hope to publish as much of the feedback as possible.

2. Have a conversation with Human Level

We are hosting 30-minute conversations with stakeholders who would like to provide their reflections and insights, share their experiences and test ideas about AI and HRDD.

To arrange a conversation, contact us at team[at]wearehumanlevel.com.

What is outside the scope of this paper

Two important dimensions of AI and human rights sit outside the scope of this paper.

First, we do not examine the human rights impacts connected to the development and use of AI itself, which increasingly form part of companies' own human rights risk profiles, including impacts on workers involved in developing, training and supporting AI systems, as well as people affected by AI-enabled products, services and decisions.

Second, we do not explore the broader implications of AI adoption for companies' business models and their resulting human rights risk profiles. As AI changes how companies operate, organise work, make decisions and create value, it will likely create new human rights risks, exacerbate existing ones or change which impacts are most severe.

Both are critical areas for HRDD. This paper has a deliberately narrower focus: the opportunities and potential pitfalls that arise when AI is used as a tool to conduct HRDD itself.

About Human Level

Human Level is an expert advisory firm advancing responsible business in practice. We specialise in business and human rights and the just transition. We partner with companies across sectors to build practical approaches that respond to today's challenges while preparing for tomorrow's expectations, turning principles and commitments into the policies, systems and decisions that shape everyday business.

This discussion paper is provided as part of our non-profit arm, Human Level Exchange, to further the discussion on meaningful HRDD. You can sign up to our weekly listserv to continue to receive updates on these topics here: wearehumanlevel.com/content-hub

POLICY DEVELOPMENT

Where AI can support, and what still needs weighing, against UNGP and CSDDD expectations for a human rights policy commitment.

UNGP / CSDDD expectation	**Opportunity:** possible AI uses	**Food for thought:** questions to consider
Set out the company's commitment to respect internationally recognised human rights. (UNGP 16)	• Propose an overall policy structure aligned with UNGP baseline expectations and practice, helping ensure no relevant sections are missed. • Compare expectations arising from soft law frameworks (UNGPs, OECD Guidelines) with binding legal requirements (e.g. CSDDD, national laws) to flag areas of overlap and divergence for legal counsel to explore further. • Generate and refine draft policy language for clarity, readability, consistency and accessibility across company documents. • Review and compare peer company policies and sector commitments to identify trends and benchmark the company's own drafting.	• Does the resulting text reflect the company's actual level of ambition and its most salient human rights issues? • Has AI-assisted legal/soft-law comparison been verified and corrected, as the case may be, by legal counsel and human rights expertise before being relied on? • Which internal conversations, alignment and reflection were compressed in producing a fast first draft, and what does this mean for the depth and ownership of the eventual commitment? • Can AI distinguish genuinely leading practice from lowest-common-denominator or outdated peer language? • Does the AI-generated text capture actual good and leading practices?
Be approved at the most senior level of the company. (UNGP 16)	• Prepare concise briefings that synthesise complex external and internal inputs (regulatory, investor, peer, internal and external stakeholder expectations) for senior leadership review. • Support scenario or stress-testing exercises that help leaders see how proposed commitments would play out in specific operational or risk contexts before sign-off.	• Does AI-prepared material bring leadership closer to the substance of what they are approving? • How is genuine leadership ownership and accountability for the policy's substance preserved as AI takes on more of the preparatory work?
Be informed by relevant expertise. (UNGP 16)	• Support informed internal and external stakeholder and expert engagement by preparing tailored materials, identifying key questions for discussion, and synthesizing relevant feedback received. • Process large volumes of external information to surface emerging issues relevant to the policy (e.g. related to civic space, climate and just transition, migration and labour trends, geopolitical instability, digital/AI governance). • Summarise the range of general expectations coming from civil society organisations, experts and trade unions.	• Where could AI create a false sense of confidence that relevant internal and external expertise has been sufficiently sought, heard and incorporated? • Where does AI-sourced input end and direct engagement with human rights experts, trade unions, worker organisations, proxies and potentially affected stakeholder begin, and is that line being maintained in practice? • Can AI reliably pick up expectations and direction-of-travel debates that are still emerging and not yet reflected in published sources?

Be communicated internally and externally. (UNGP 16)	• Support translation of the policy into multiple languages and adapt communications for different audiences, geographies and literacy levels. • Help draft internal communications, training materials and leadership messaging that reinforce the policy's substance across the business.	• Does translation or simplification preserve the nuance and intent of the original commitments, or risk losing meaning along the way? • Is the translation or simplification fit for purpose for the audience, geography and context in which it will be used? • Is AI-supported communication paired with genuine internal engagement and dialogue, leading to the conversations needed to build real understanding and ownership?
Integrate due diligence into policies and risk management systems (CSDDD, Article 7)	• Monitor changes in legal requirements, stakeholder expectations, industry standards and emerging general human rights risks over time, flagging when the policy may need review. • Review existing policies, governance documents and operational procedures to identify where human rights expectations are not yet integrated or are inconsistent across systems.	• How will AI-identified gaps or review triggers be verified? • Does treating policy-risk-management integration as a monitoring exercise risk overlooking the deeper organisational and governance change that Article 7 is intended to drive?

POLICY EMBEDDING

Where AI can support, and what still needs weighing, to embed respect for human rights throughout the business.

UNGP / CSDDD expectation	**Opportunity:** possible AI uses	**Food for thought:** questions to consider
Internal alignment (UNGP 16 & 19)	• Prepare internal communications, speeches, briefing notes, leadership messages and engagement materials that support the 'tone from the top' on human rights. • Help employees access relevant policies, guidance, escalation pathways and human rights resources through AI-enabled internal tools. • Tailor implementation guidance to different business functions – e.g. procurement, HR, security, sales, product development – so that each understands what the policy means for its own day-to-day decisions. • Review existing policies, governance documents, operational procedures and internal controls to flag where human rights expectations are missing or inconsistent across systems.	• Does AI-drafted 'tone from the top' messaging read as an authentic leadership voice and commitment? • When function-specific guidance is generated at scale, who checks that it reflects the actual operational realities and risk profile of that function, rather than generic good practice? • Does easier access to policies and guidance change how much employees genuinely engage with the substance?

Training and capacity-building (UNGP 16 & 19)	• Prepare training materials, FAQs, case studies, scenario exercises, presentations and awareness materials to help employees and business partners understand what commitments mean in practice. • Tailor training content for different audiences, geographies, languages, literacy levels and operational contexts.	• Does scaling training through AI trade depth and dialogue for reach and volume? • How do we know whether AI-personalised training is actually landing with different audiences? • Where in the training experience is there still room for questions, discussion and challenge?
Leadership commitment (UNGP 16 & 19)	• Help leaders prepare for engagement moments such as town halls, board briefings, and supplier summits by synthesising relevant human rights context, risk data and stakeholder feedback into concise briefings. • Support scenario exercises that help leaders see what specific commitments would mean for real operational and commercial decisions.	• Does AI-prepared briefing material bring leaders closer to the substance of the issues? • How is visible, authentic leadership commitment preserved as more of the supporting material becomes AI-generated?
Accountability and incentive structures (UNGP 16 & 19)	• Help design or refine KPIs, incentive structures and performance metrics that reflect human rights responsibilities across relevant roles and functions. • Help identify where existing KPIs may conflict with human rights objectives, and where human rights KPIs themselves could create unintended consequences or perverse incentives. • Support tracking of who owns which human rights-related actions, decisions and follow-ups across the business, and flag where items are stalling. • Flag internal incentive structures, KPIs and performance metrics for buyers, sourcing teams, and country managers that may reward behaviour that conflicts with the human rights policy commitment, such as cost-per-unit or speed-to-market targets with no reference to labour or sourcing risk. • Help map out where accountability for human rights currently sits across similar companies' governance structures (e.g. board committee, general counsel, chief sustainability officer, dedicated human rights lead), surfacing options a company could consider adapting to its own structure and maturity.	• Can AI meaningfully weigh trade-offs between commercial incentives and human rights considerations? • Can AI genuinely identify where KPIs may undermine human rights objectives or create unintended consequences, without understanding the human realities and behaviours behind the data? • Where AI helps surface ownership and follow-through gaps, who is accountable for actually acting on what it flags? • Can AI recommend an ownership structure that is genuinely tailored to this company's size, culture and risk profile, rather than one that is simply common among peers?

Align contracts with the human rights policy commitment. (UNGP 16 & 19; CSDDD, Article 7)	• Review supplier, joint venture, distributor and other commercial contracts at scale to determine whether human rights clauses are present and how consistent they are with the policy commitment and the principles of responsible contracting. • Review commercial contracts at scale to determine whether human rights respect may be undercut by other contractual terms, for example through pricing, lead times, liability allocation, termination provisions or other commercial requirements. • Help draft or refine responsible contracting clauses, drawing on evolving good practice and focusing on provisions that support meaningful changes in practice. • Help model how a proposed change to commercial terms such as a shorter lead time, a lower price point, or a volume commitment might interact with known human rights risk factors in the relevant part of the value chain, before the terms are finalised. • Support monitoring of whether responsible contracting commitments are implemented in practice, including whether parties cooperate to address impacts, share responsibility for corrective action and remedy, and use escalation or termination responsibly.	• Does a contract containing a human rights clause become evidence of embedding in itself, where the underlying commercial relationship and incentives remain unchanged? • Can AI effectively identify where a contract's commercial terms work against its own human rights clauses? • If an incentive structure rewards a buyer for driving down costs and AI flags a related human rights risk, what actually happens? • Where AI suggests a responsible contracting clause, is it aligned with current good practice, and has it been checked against current legal thinking on what actually works in practice?
Integration into day-to-day business activities and relationships (UNGP 16 & 19)	• Identify where human rights considerations need to be built into existing decision-making processes: supplier onboarding, M&A, sourcing decisions, product development, investment decisions, market entry, contract approval, etc. • Support development of supplier scorecards or evaluation frameworks that incorporate human rights criteria alongside cost, quality and delivery. • Identify patterns in purchasing practices, including for example related to delivery schedules, pricing demands, order changes and contract volatility, that may be contributing to human rights risks. • Help prepare supplier guidance, onboarding materials and communication resources that support understanding of the company's expectations across the value chain.	• Does adding human rights criteria to scorecards and contracts actually change supplier behaviour and outcomes? • Could AI create a false sense that human rights are "integrated" because they appear in workflows, scorecards or approval processes, without changing underlying incentives and behaviours? • Does the data available to AI capture the commercial pressures experienced by suppliers and workers, or only what is visible through company systems?
Integrate due diligence into policies and risk management systems (CSDDD, Article 7)	• Help identify where existing risk management systems, governance frameworks and internal controls are not yet capturing human rights. • Support ongoing monitoring of how due diligence findings feed back into policies, systems and decision-making, and flag where integration gaps persist over time.	• Does treating this as a systems or technical integration exercise risk missing the deeper organisational and cultural embedding that Article 7 is intended to drive? • Who owns the follow-through when AI flags a gap between due diligence findings and existing risk management systems?

RISK ASSESSMENT AND PRIORITISATION

Where AI can support, and what still needs weighing, to identify, assess and prioritise human rights risks and impacts.

UNGP / CSDDD expectation	**Opportunity:** possible AI uses	**Food for thought:** questions to consider
Identify and assess actual and potential adverse human rights impacts connected to own operations and the value chain (UNGP 17 & 18)	• Process and analyse large volumes of internal information relevant to human rights risks, such as audit findings, grievance data, supplier assessments, operational incidents and worker feedback. • Process and analyse large volumes of external information relevant to human rights risks, such as media reporting, civil society and trade union publications, academic research and social media discourse, including material previously less accessible due to language barriers. • Identify recurring patterns and systemic issues across complaints, incidents, supplier behaviour or operational practices that may point to deeper systemic risks. • Support horizon scanning, trend analysis and predictive assessment across factors such as conflict escalation, climate impacts, migration flows, labour shortages, commodity price volatility, political instability, supply chain disruption, recruitment patterns or sudden sourcing changes. • Identify where low levels of reported grievances or audit findings may not indicate low risk, for instance where fear of retaliation, restricted worker voice, weak civic space or limited access to remedy are more likely explanations. • Highlight links between business practices, like sourcing decisions, commercial incentives, production timelines, or labour models, and heightened human rights risks.	• Can AI reliably interpret how a human rights risk is experienced by the people affected? • How does AI account for local context, power dynamics and informal structures that may not appear in company data or public reporting? • Where data is missing or thin, does that get treated as low risk, or as a signal that deeper human inquiry is needed?
Prioritise risks and impacts based on severity and likelihood (UNGP 17 & 18)	• Support a structured scoring or heat-mapping approach across scale, scope and remediability (the three components of severity) to make comparison more consistent across geographies, operations, suppliers and workforce groups. • Identify where groups associated with heightened vulnerability, such as migrant workers, women, children, or Indigenous Peoples, are present across operations and the value chain, feeding into severity assessment. • Identify compounding vulnerabilities where multiple risk factors intersect, for example migrant workers in conflict-affected areas, or labour shortages alongside extensive subcontracting in weak governance settings.	• Whose judgment sets the weightings behind an AI-supported scoring model, and how are affected stakeholders' perspectives reflected in that judgment? • Does an AI-supported scoring model still allow one severity factor, such as scale, to drive prioritisation on its own, without high scope or low remediability? • Does prioritisation by AI tend to favour risks with more available data over severe but less-documented risks?

	• Support analysis of likelihood factors at scale, such as reliance on third-party recruitment agencies, complex subcontracting chains, short-term supplier relationships, weak rule of law, conflict or restricted civic freedoms.	
Focus on risks to people, rather than risks to the business (UNGP 17 & 18)	• Help structure risk registers, dashboards and reporting categories around impacts on and risks to people rather than business risk categories such as reputational or legal exposure. • Flag where existing risk frameworks or data structures are oriented toward risk to the business, prompting a review of how human rights-related risks are categorised and framed.	• Do the metrics AI relies on actually measure risk to people, or are they proxies for risk to the business such as cost, disruption, and reputational exposure? • Does the way AI visualises or ranks risk data unintentionally reintroduce a business-risk lens, by surfacing what is most visible to the company rather than what matters most to the people affected/ potentially affected?
Draw on engagement with potentially affected and other relevant stakeholders, as well as appropriate human rights expertise (UNGP 17, 18 & 24)	• Help map relevant stakeholders and identify vulnerable groups or categories of rights-holders that appear underrepresented in existing assessment processes. • Prepare stakeholder mapping exercises, briefing materials and consultation questions ahead of engagement. • Analyse and summarise large volumes of stakeholder feedback, consultation outcomes, grievance data and worker voice information at scale. • Support multilingual engagement by translating and adapting materials for different stakeholder groups and literacy levels.	• Can AI-supported stakeholder mapping identify who is missing from the process, or only who is already represented in existing data? • Does AI-assisted analysis of stakeholder feedback capture what is left unsaid, such as concerns withheld out of fear or distrust? • Where AI supports engagement preparation, is direct human engagement with affected stakeholders still happening, or does the preparation start to stand in for it?
Identify and assess adverse human rights and environmental impacts across own operations, subsidiaries, and chains of activities (CSDDD, Articles 8 & 9)	• Extend risk assessment to cover environmental impacts alongside human rights impacts, consistent with CSDDD's broader scope. • Map and analyse chains of activities more systematically, including indirect business relationships. • Identify where operations, subsidiaries or chains of activities are not yet covered by existing risk assessment processes.	• Can AI reliably distinguish between direct operations, subsidiaries and the wider chain of activities, where CSDDD obligations and company leverage differ significantly? • Does AI have the necessary data to undertake the mapping and extend the risk assessment? • Does AI-supported mapping of chains of activities create a sense of complete visibility over indirect relationships?

STAKEHOLDER ENGAGEMENT

Where AI can support, and what still needs weighing, to engage meaningfully with potentially affected people and other relevant stakeholders.

UNGP / CSDDD expectation	**Opportunity:** possible AI uses	**Food for thought:** questions to consider
Engage the right people: meaningfully consult potentially affected people and relevant stakeholders. (UNGP 18 & 24; CSDDD, Article 13)	• Map relevant stakeholders and identify vulnerable stakeholder groups or categories of rights-holders that appear underrepresented or insufficiently considered in existing assessment processes. • Cross-reference operational and value chain data like locations, supplier lists, or workforce composition against known risk factors to flag stakeholder groups that may otherwise be overlooked. • Identify existing engagement channels, forums or intermediaries for example unions, NGOs, community groups that are already connected to a given stakeholder group, to inform who to approach and how.	• Can AI-supported stakeholder mapping identify who is missing from the process, or only who is already visible in existing data and systems? • Does reliance on digital or documented sources systematically miss stakeholders who are informal, undocumented or offline? • How does AI account for the lived experience and human impacts that only become visible through direct engagement, rather than through available data?
Give stakeholders enough information to participate effectively. (UNGP 18; CSDDD, Article 13(3))	• Prepare accessible briefing materials, FAQs and background information that explain the purpose, scope and process of an engagement in advance. • Tailor explanatory materials to what a specific stakeholder group needs to know to meaningfully participate, rather than sending generic information.	• Does the information AI helps prepare give stakeholders a genuine basis to participate, or does it lean toward what's easiest for the company to produce? • Who checks that stakeholders have actually understood the information well enough to participate meaningfully, rather than assuming that sharing it is sufficient?
Engage at meaningful points in due diligence. (UNGP 18; CSDDD, Article 13(4))	• Help identify at what point in a risk assessment, action-planning or remediation process stakeholder input would be most useful, based on where decisions are still open. • Flag upcoming decisions, sourcing changes or planned actions across the business early enough for engagement to be planned before decisions are effectively locked in. • Support scheduling, logistics and follow-up coordination across multiple engagement touchpoints over time.	• Does AI help engagement happen earlier, while decisions are still open, or does it mainly support engagement after decisions have already been made? • Who decides which points in the due diligence process count as genuinely meaningful for engagement, rather than leaving that judgment to an AI-generated timeline?

Use credible alternatives where direct engagement isn't reasonably possible. (UNGP 18; CSDDD, Article 13(4))	• Help identify credible secondary sources – trusted civil society, trade unions, academics, experts – that can stand in where direct engagement isn't feasible. • Support triangulation across multiple expert sources to help assess how credible or representative a given alternative source is.	• Who determines that direct engagement is genuinely not reasonably possible, rather than simply more difficult or costly, before turning to an alternative? • Does an AI-suggested secondary source reflect the most relevant expertise, or simply the voices that are most visible online?
Draw on affected stakeholders' feedback when tracking whether responses are working, and communicate how impacts are being addressed. (UNGP 20 & 21; CSDDD, Article 13(4)(e))	• Identify recurring stakeholder concerns and trends by analysing common themes, patterns or emerging issues raised by workers, communities, trade unions, civil society organisations or other relevant stakeholders. • Synthesise stakeholder feedback collected over time to show whether concerns are recurring, easing or escalating. • Support drafting of feedback loops or updates that explain to stakeholders how their input has, or hasn't, shaped the company's response.	• Does an AI-identified improvement in sentiment or a drop in complaints reflect real change, or could it reflect fear, fatigue, or declining trust in the process? • Where stakeholder concerns don't lead to visible action, how is the feedback loop closed with those stakeholders in a way that is meaningful and appropriate for them?
Take account of barriers such as language. (UNGP 18; CSDDD, Article 13(5))	• Support translation and communication across multiple languages, including smaller or less commonly supported languages. • Simplify or tailor communications, explanatory materials or engagement documents for different stakeholder groups, geographies, contexts and literacy levels. • Adapt materials into different formats such as audio, visual, or simplified text to reach stakeholders with different literacy levels or access needs.	• Does AI-assisted translation and simplification preserve the tone, nuance and intent of the original message, particularly for culturally or legally sensitive content? Is it fit for purpose for the audience, geography and context in which it will be used? • Who reviews AI-translated or simplified materials before they reach stakeholders, to catch errors or loss of meaning that could affect trust? • Are there language or dialect groups within the value chain that fall outside what AI translation tools currently support well?
Make engagement genuinely accessible and safe. (UNGP 18; CSDDD, Article 13(5))	• Support multilingual, multi-channel and multi-format engagement options (in person, phone, written, digital) so stakeholders can participate in ways that suit their circumstances. • Help identify potential safety or retaliation risks associated with specific engagement formats or locations, based on known context factors, so these can be planned around.	• As AI is used, how is human engagement preserved where trust, empathy and psychological safety are important for people to feel able to speak openly? • Does the involvement of AI in collecting or analysing stakeholder input create new concerns about privacy, surveillance or data misuse that could discourage participation? • How is fear, hesitation or reluctance to speak interpreted when it isn't captured as data at all?

ACTION AND REMEDIATION

Where AI can support, and what still needs weighing, in acting on human rights findings and remediating impacts.

UNGP / CSDDD expectation	**Opportunity:** possible AI uses	**Food for thought:** questions to consider
Integrate findings from human rights impact assessments across relevant internal functions and processes, and take appropriate action in response. (UNGP 19)	• Identify links between operational practices and human rights impacts (purchasing practices, pricing pressures, production targets, recruitment models, incentive structures) so action can target root causes rather than symptoms alone. • Help route assessment findings to the specific business where the root causes and levers for action actually sit. • Support integration of findings into existing decision-making forums such as sourcing reviews, budget cycles, or supplier scorecards.	• Does routing findings to specific functions via AI actually change what those functions decide? • Can AI reliably distinguish a root cause from a symptom, or does it tend to flag whichever pattern is most visible in the data available to it? • Who is accountable for ensuring a finding is actually acted on once AI has routed or flagged it?
Where a company causes or may cause an adverse impact, take steps to cease or prevent it. (UNGP 19) **Take appropriate measures to identify, prevent, mitigate, bring to an end, and minimise adverse human rights and environmental impacts arising from own operations, subsidiaries, and chains of activities.** (CSDDD, Articles 10 & 11)	• Help identify where the company's own operations, decisions or practices are the most likely direct source of a risk or impact, to distinguish 'cause' situations from other situations. • Support scenario testing of what ceasing or changing a specific internal action would mean operationally, before a decision is made.	• Can AI reliably tell the difference between a company causing an impact and other forms of involvement? • Where AI recommends ceasing a practice, does that recommendation account for any potential operational or livelihood consequences of stopping it abruptly?

Where a company contributes or may contribute to an impact, cease or prevent its contribution and use leverage to mitigate any remaining impact. (UNGP 19; and CSDDD, Articles 10 & 11)	• Help distinguish the company's own contribution to an impact from the conduct of the business partner responsible for the remaining harm, to clarify what the company needs to change itself versus where it needs to use leverage. • Assess the robustness of proposed mitigation measures by identifying potential gaps, dependencies, implementation risks, unintended consequences or missing information. • Support scenario analysis to explore how mitigation measures may perform under different operational, commercial, geopolitical or environmental conditions, flagging where additional safeguards may be needed. • Aggregate and analyse how peer companies, industry initiatives and multistakeholder initiatives are responding to specific risks, alongside recommendations from civil society, trade unions, workers, Indigenous Peoples, UN bodies and OECD National Contact Points. • Identify where expectations on mitigation and prevention are evolving in emerging standards, investor expectations, regulatory developments, sector commitments, and civil society advocacy. • Connect root causes to specific prevention or mitigation actions by linking purchasing practices, pricing pressures or recruitment models to the impacts they are associated with, across operations, subsidiaries and the wider chain of activities. • See row below	• Can AI reliably separate the company's own contribution to a harm from a business partner's independent conduct and thus assign accountability? • Does AI-supported stress-testing of mitigation measures capture how a measure will actually play out for affected people, or only whether it looks internally consistent on paper? • Could AI favour mitigation measures that are easier to quantify or implement over those most likely to prevent or reduce harm to people? • Does treating widespread peer practice as evidence of effectiveness risk reproducing weak or lowest-common-denominator responses? • See row below
Where an impact is directly linked to the company through a business relationship, use or increase leverage to prevent or mitigate it. (UNGP 19; and CSDDD, Articles 10 & 11)	• See row above • Help identify the full range of leverage points available to the company (e.g., through contracts, the nature of the commercial relationship, purchasing power, industry collaboration or financing arrangements) and suggest ways to build and exercise leverage effectively in light of the specific relationship and context. • Identify where collaboration with industry peers, multistakeholder initiatives, trade unions, civil society organisations, governments or investors could strengthen leverage on systemic risks, and surface existing collaboration opportunities.	• See row above • Can AI fully understand the relationships, power dynamics and informal sources of influence that determine where leverage actually exists and how it can be exercised? • Does AI-identified leverage reflect what's realistically exercisable in the relationship, or mainly what's visible in contracts and formal documentation? • Could relying on AI's assessment of leverage cause a company to underestimate leverage that comes from long-standing relationships or informal influence not captured in written records?

15 **Where a company has caused or contributed to an adverse human rights impact, provide for or cooperate in remediation through legitimate processes.** (UNGP 19 & 22) **Establish a notification mechanism and complaints procedure, and provide remediation where the company has caused or jointly caused an adverse impact.** (CSDDD, Articles 14 & 16)	• Identify patterns of recurring impacts across suppliers, geographies or business units, supporting more systemic remediation responses rather than isolated case-by-case handling. • Support case management by organising and summarising documentation, tracking timelines and follow-up actions, and triaging and routing incoming cases while flagging urgent or high-severity issues for human attention. • Help identify barriers to effective remedy such as low grievance uptake, repeated unresolved complaints, delayed responses, disengagement from processes. • Make notification channels and complaints mechanisms more accessible across languages, literacy levels, communication formats and geographies. • Track whether complaints are being resolved within reasonable timeframes and flag recurring or unresolved patterns that may indicate a systemic issue or a weakness in the mechanism itself.	• Does faster case processing through AI reflect more effective remedy for the people affected? • Can AI capture what affected people themselves consider a fair, appropriate or restorative outcome, or does 'remediation' risk being defined by what's easiest to document as closed? • Where AI is involved in handling sensitive complaint or remediation data, what new privacy or retaliation risks does that introduce for complainants? • Does AI-supported triage speed up attention to genuinely urgent and severe cases, or could it deprioritise complaints that don't fit familiar patterns or categories? • Are complainants told, and do they understand, when and how AI is involved in handling their notification or complaint? • Would the people most likely to need this mechanism, particularly the most vulnerable, trust and feel safe using a channel they know involves AI processing of their information?

TRACKING

Where AI can support, and what still needs weighing, in tracking whether responses to human rights risks and impacts are actually working.

UNGP / CSDDD expectation	**Opportunity:** possible AI uses	**Food for thought:** questions to consider
Track the effectiveness of responses to human rights risks and impacts, to assess whether impacts are being properly prevented and mitigated. (UNGP 20)	• Help monitor implementation and progress against action plans for salient human rights issues and related KPIs, flagging delays, gaps or emerging challenges. • Help monitor supplier progress on corrective actions, flag overdue items, and track whether issues reappear after actions are closed. • Detect trends and recurring themes over time, for example repeated wage complaints, recurring excessive overtime, unresolved grievances, and deteriorating worker sentiment, that may signal existing responses aren't working. • Support more dynamic board and management oversight through dashboards that track trends and flag where mitigation efforts may not be delivering intended outcomes.	• Does improvement in a dashboard or KPI reflect meaningful improvement in outcomes for affected people? • Are the people most at risk adequately represented in the data being used to track progress? • When an indicator improves, such as fewer complaints or flagged issues, how do we know this reflects reduced impact rather than under-reporting, reduced access or other factors? • Are indicators remaining consistent enough over time to show genuine progress, rather than changes in definitions, methodology or data availability? • Is AI able to identify meaningful changes and emerging concerns that predefined KPIs may miss, including qualitative signals and unexpected outcomes?

16 **Use qualitative and quantitative indicators, drawing on internal and external sources of feedback – particularly from affected stakeholders.** (UNGP 20)	• Help compile inputs into key indicators by reviewing large datasets from grievance mechanisms, audits, corrective action plans, worker surveys, turnover data, health and safety reports and supplier assessments. • Help analyse large volumes of qualitative stakeholder feedback to identify patterns, changes in experience and signals of whether conditions are improving. • Help prepare for targeted stakeholder engagement to test whether KPIs reflect affected stakeholders' experiences and identify gaps or outcomes that existing indicators may be missing.	• Whose feedback is represented in the data AI analyses, and whose experiences may be missing or underrepresented? • Can AI-supported analysis capture what people may choose not to say because of fear, distrust, power dynamics or previous experience? • Could analysis of existing stakeholder feedback become a substitute for the direct, ongoing engagement needed to understand whether conditions are actually changing? • Are quantitative indicators being given greater weight simply because they are easier for AI to analyse, at the expense of qualitative stakeholder experience?
Carry out periodic assessments to monitor the adequacy and effectiveness of measures taken, using qualitative and quantitative indicators where appropriate, and after a significant change occurs (at least every five years). (CSDDD, Article 15)	• Help schedule and structure periodic effectiveness assessments, and flag changes in operations, business relationships or external context that may warrant an earlier review. • Help assemble and synthesise relevant qualitative and quantitative indicators from across existing systems, reducing the manual work involved in periodic assessments. • Help compare performance over time and identify trends, inconsistencies or areas where measures do not appear to be producing the intended outcomes. • Help identify gaps in the evidence base and areas requiring further investigation or stakeholder engagement before conclusions on effectiveness are reached.	• Does a periodic assessment built mainly on existing indicators still capture emerging or newly significant changes? • How can changes in operations, business relationships or external context be visible in the data available to AI? • Who determines whether a change is significant enough to warrant reassessment, and is that judgement being delegated too readily to AI?

COMMUNICATING

Where AI can support, and what still needs weighing, in communicating externally on human rights due diligence.

UNGP / CSDDD expectation	**Opportunity:** possible AI uses	**Food for thought:** questions to consider
Communicate externally how human rights impacts are addressed, with formal reporting expected where severe human rights risks exist. (UNGP 21)	• Organise and synthesise due diligence-related information into coherent reporting outputs. • Help align disclosures across different frameworks and reporting requirements and quickly update and refine disclosures as data and inputs evolve.	• Could AI make it too easy to produce polished, reassuring disclosures quickly, including where the underlying due diligence work behind them is still weak? • Does AI-assisted drafting soften language around severe impacts, controversies or unresolved issues in ways that lose the nuance needed to understand what's actually happening? • Where severe risks exist, does the resulting report reflect that severity, or does confident, well-structured AI-generated language make things look more resolved than they are?
Communicate in a form and frequency that reflects the company's human rights risks. (UNGP 21)	• Support more timely and dynamic communication beyond the annual reporting cycle through regular updates, public dashboards, or targeted stakeholder communications. • Help translate complex legal or technical reporting into clearer, more accessible language for different audiences.	• Does more frequent AI-assisted communication reflect genuinely new developments, or does it create an appearance of dynamism without new meaningful substance behind it? • Does standardised, AI-generated language risk understating the company's risks and trade-offs in the process of making reporting more accessible?
Provide sufficient information to evaluate the adequacy of the company's response. (UNGP 21)	• Help connect disclosure more directly to stakeholder feedback, underlying operational challenges and remediation outcomes, moving communication toward demonstrating implementation and effectiveness rather than intent.	• Does an AI-assisted report give readers sufficient information to evaluate whether the company's response is adequate? • Whose voice is present in the final disclosure, and whose lived experience gets filtered out as feedback is synthesised into a polished narrative? • Who decides what information stakeholders need to evaluate the company's human rights approach, given that AI can help prepare disclosures but may not be well placed to determine what is most important to include?

Where not otherwise subject to sustainability reporting requirements, publish an annual statement on the company's website covering the matters addressed by the Directive. (CSDDD, Article 20)	• For companies not otherwise required to report under sustainability reporting rules, help draft and structure the annual due diligence statement, keeping it aligned with what the Directive actually requires. • For companies required to report under existing sustainability reporting rules, assess potential gaps between CSDDD expectations, and CSRD reporting. • Help keep the statement up to date as due diligence processes, risks or actions change from year to year.	• Does AI help ensure the substantive disclosure requirements are met, or primarily support its format and presentation? • Who checks that an AI-drafted statement doesn't simply repeat prior years' language in areas where processes have genuinely changed?
Report on sustainability matters, including human rights and environmental due diligence, within the management report, covering governance, strategy, impact/risk/opportunity management, policies, actions, targets and metrics, subject to double materiality assessment and third-party assurance. (CSRD & ESRS)	• Help structure disclosures against the specific ESRS requirements relevant to human rights. • Help ensure consistency and traceability between the due diligence process described in the report and the underlying data, given that CSRD disclosures are subject to third-party assurance and digital tagging. • Help identify where a company's existing due diligence information doesn't yet map cleanly onto required ESRS data points, ahead of assurance.	• If disclosures are going to be assured by a third party, does AI-generated content create a stronger paper trail of what the company actually did, or a stronger paper trail of what it says it did? • Where a required data point calls for a metric the company doesn't actually track well, does AI help surface that gap honestly?

Closing reflections

Across our discussions on AI and HRDD, eight themes have surfaced consistently regardless of which stage of the due diligence process we are considering.

1. Start with HRDD, not with AI capabilities

The starting point should be what effective HRDD requires, and then where and how AI can help deliver it – rather than starting with what AI is capable of doing.

Companies already face significant challenges in conducting meaningful HRDD, from fragmented information and limited supply chain visibility to identifying emerging risks, processing large volumes of information and enabling meaningful stakeholder engagement. AI may be well suited to addressing some of these challenges, and much less suited to others.

Starting with the objectives and expectations of HRDD provides the anchor. It helps companies determine where AI can genuinely strengthen the process, what role it should play, and where its use could inadvertently weaken the quality or effectiveness of due diligence.

2. AI can analyse what is visible; HRDD needs to also find what is invisible

AI is fundamentally dependent on the information available to it. But in HRDD, some of the most important information may be precisely what is not captured in existing data.

Severe human rights impacts often occur in contexts characterised by informality, fear, under-reporting and unequal power. Those facing the greatest risks – including informal and migrant workers, Indigenous Peoples, women, human rights defenders, people with disabilities and others facing heightened vulnerability – may also be among the least visible in company systems and datasets.

Existing information may therefore be incomplete, outdated or unrepresentative, and historical data may reproduce existing assumptions, inequalities and biases. AI can identify patterns within this information at a scale humans cannot – but greater analytical sophistication does not necessarily produce a more complete picture. It can also reinforce existing blind spots. This makes understanding what is missing, whose experiences are not represented, and where additional information or direct stakeholder engagement is needed fundamental to the effective use of AI in HRDD.

3. Meaningful stakeholder engagement becomes more important, not less

At its heart, HRDD is about understanding what is happening to people and how they experience impacts on their human rights. That understanding cannot come from data alone.

Whether someone feels able to speak openly about an impact can depend on trust, fear, power dynamics, cultural context and risks of retaliation. What remains unsaid may be as important as what is captured in a dataset.

AI can help companies analyse information from stakeholder engagement, identify patterns and prepare for further conversations. But it cannot substitute for the relationships and trust through which impacts are surfaced and understood.

Indeed, the greater the role AI plays in analysing HRDD information, the more important meaningful stakeholder engagement becomes – both to bring perspectives into the process that would otherwise remain invisible, and to test, challenge and add context to what AI produces.

4. AI should augment human judgment not replace it

Across almost every stage of HRDD, the same dividing line emerges: the difference between informing a decision and making one.

AI might summarise information, assess severity, suggest priorities, evaluate leverage, propose indicators or draft disclosures. But there is a risk of decision creep: summarising becomes assessing; assessing becomes recommending; and recommending begins, in practice, to determine what happens next.

The more authoritative or sophisticated an AI output appears, the easier it may be for that distinction to blur.

Companies therefore need to be deliberate about where AI's role ends and human judgment begins. Who interprets the output? Who challenges it? Who makes the ultimate decision? And who remains accountable for its consequences?

AI can significantly expand the information and analysis available to decision-makers. It should strengthen human judgment, not substitute for it or obscure responsibility for the decisions ultimately taken.

5. The capability around AI matters as much as the capability of AI

The effective use of AI in HRDD depends not only on what the technology can do, but on the capability and judgment of the people using it.

Practitioners need more than the technical skills to prompt and operate AI tools. They need to be able to interrogate outputs, recognise what may be missing or misleading, identify where further information or stakeholder engagement is needed, and know when an AI-generated output should be challenged or not relied upon.

This requires capability-building that is grounded in the objectives and principles of HRDD, not simply in how to use AI. As AI becomes more integrated into due diligence, practitioners will need to understand both its potential and its limitations, while retaining the confidence and authority to exercise their own judgment.

The goal should not simply be to make practitioners better at using AI, but better at questioning it – and at knowing when human judgment, expertise and engagement need to take over.

6. Leverage with AI service providers will become increasingly important

As companies integrate AI into HRDD, the quality of their due diligence will increasingly depend on systems they did not design and cannot fully see, understand or control.

Many of the questions raised throughout this paper – about data sources, representativeness, bias, explainability, confidentiality, human oversight and the ability to challenge outputs – cannot be addressed by the company using the tool alone. They also require answers, information and, potentially, changes from AI service providers.

This makes leverage with AI providers an increasingly important part of the conversation.

Companies can consider these issues when selecting providers, conducting procurement and due diligence, negotiating contracts, setting expectations and reviewing performance. They can ask what data systems rely on, how limitations and biases are identified, what transparency is available around outputs, what happens to information entered into the system, and what mechanisms exist to raise concerns or seek changes.

But individual companies may have limited leverage, particularly with large technology providers. Collective leverage is therefore likely to become increasingly important: companies, investors, industry initiatives and others aligning around common expectations of AI providers could help establish clearer standards for what responsible AI used in HRDD needs to enable.

7. Measure whether AI is actually making HRDD better

AI offers potential to make HRDD faster, more scalable and more efficient – from analysing large datasets and identifying patterns to triaging grievances, monitoring indicators and supporting disclosure. But efficiency is not the ultimate measure of successful HRDD. Effectiveness is.

Companies therefore need to look beyond whether AI saves time or processes more information and ask whether it is actually improving the quality and outcomes of due diligence. Is it helping identify impacts that were previously missed? Are emerging concerns being identified earlier? Is the company gaining a better understanding of severity and context? Is stakeholder engagement becoming more meaningful? Are actions and remedy becoming more effective?

There is also an important test of what happens to the time AI saves. Does greater efficiency create more space for practitioners to engage with affected stakeholders, exercise judgment and focus on action – or does it simply enable companies to process more information and produce more outputs?

Ultimately, the measure of AI's value to HRDD should be not how much more due diligence activity it enables, but whether it contributes to more effective due diligence and better outcomes for people.

8. Using AI for HRDD also brings AI into the scope of HRDD

There is an important connection between using AI to conduct HRDD and conducting HRDD on the company's use of AI itself.

As companies procure and integrate AI tools into their due diligence processes, they introduce new business relationships, suppliers and potential human rights impacts. These may include impacts connected to the workers who develop, train and maintain AI systems; data sourcing and privacy; discrimination and bias; surveillance; environmental impacts associated with AI infrastructure; and impacts on people affected by AI-supported decisions.

This means that introducing AI into HRDD does not simply change how due diligence is conducted – it also changes what needs to be subject to due diligence.

Companies therefore need to consider the AI systems and providers they rely on within their broader HRDD processes: identifying and assessing associated impacts, taking appropriate action, using leverage with providers, tracking effectiveness and addressing impacts where relevant.

There are still a number of significant open questions about how responsibility and leverage can operate across complex AI value chains. But the underlying principle remains: as AI becomes a tool for conducting HRDD, the company's use of AI – and the impacts connected to it – also become part of a company's HRDD picture.

We look forward to hearing your perspectives, experiences and questions – and to continuing to shape this conversation together.