

- 
-  Sustainable
 -  Finance
 -  Cluster

POSITION PAPER

AI in ESG Management

Adoption Maturity, Use Cases
and Responsible Governance
in Banking

September 2026

Table of Contents

Executive Summary	3
Introduction	5
1. AI Adoption Maturity	6
Context	6
Implementation Maturity	7
Drivers, Barriers and Challenges	8
Recommendations for Banks	10
2. Use Cases	11
Context	11
Selection Criteria, Adoption and Value	11
Overview of use cases and case studies	16
Helaba Case Study: ESG analysis in the credit process	19
Novata Case Study: Reputational risk management	21
Commerzbank Case Study: Screening of greenwashing risks	23
Recommendations for Banks	25
3. Responsible & Sustainable AI Governance	26
Context	26
Responsible AI	27
Sustainability of AI	28
Recommendations for Banks	30
Outlook: From Scaling to Agentic AI	31
 References	 32
Appendix	34
Approach of the Sustainable Finance Cluster Initiative	39
Acknowledgements	40
Imprint	41

Executive Summary

Environmental, social and governance (ESG) management has become increasingly data-intensive and regulated, and banks are turning to artificial intelligence to keep up the pace. Much of the ESG workload is slow, manual work, from gathering data across scattered sources to reading and summarising sustainability reports and preparing disclosures. This is where AI helps first. General-purpose AI is already in use across banks' functions, so in the context of ESG the question is no longer whether to use it, but how to move from isolated pilots to deliberate and responsible scaling. Drawing on workshops, case studies and a survey of the banks in the Cluster Initiative, this paper describes the current status of AI use for ESG management in banking. It also identifies barriers and levers that could accelerate adoption. The overall picture is a sector that has committed to the topic but is still early in practice. The value is real but concentrated in efficiency gains, and it will come less from more advanced technology than from how deliberately banks scale, govern and account for the AI they use.



IMAGE SOURCE: [Thurston H/peopleimages.com_600679051](https://www.gettyimages.com/detail/stock-photo/600679051) - stock.adobe.com

This paper identifies five key takeaways:

1. FROM PILOTS TO SCALE – Reuse, Not Repetition

Banks are actively exploring AI for ESG, but adoption remains uneven and rarely scaled beyond individual use cases. General-purpose chatbots are used everywhere, while deeper and agentic uses are rare. Banks should scale proven use cases deliberately, while conceptualising and piloting selected agentic AI applications within clear guardrails. Banks should treat buy-and-build as a portfolio decision, buying to enter quickly while developing reusable capabilities on a shared foundation, so that each use case adds to a base that scales rather than standing alone.

2. DURABILITY OVER SPEED – Choose Well and Measure Value

Banks often prioritise technically feasible use cases, but durable scaling also depends on data availability, regulatory fit, user adoption and measurable value. Use cases should therefore be selected not only for what can be built quickly, but for what can be governed, adopted and shown to improve ESG work. Value measurement should support adoption by capturing quality, reduced manual burden, redeployed capacity and ESG-specific outcomes, not only time or cost savings.

3. PUT A CLEAR OWNER ON IT – Governance and the EU AI Act

As AI enters high-stakes processes such as credit and disclosure, governance becomes central, yet no single ownership model has emerged. Banks should give AI in ESG a clear owner whose remit includes model risk, get ahead of the EU AI Act's high-risk obligations, and make responsible-AI safeguards such as source-checking, human oversight and documentation standard rather than ad-hoc.

4. MANAGE THE DEPENDENCY – Digital Sovereignty as a Risk

Models and hardware come from a few mostly non-European providers, a concentration the banks do not yet rate as a pressing barrier but which is a potential operational risk, as the mid-2026 access restrictions have demonstrated. Withdrawing is not realistic, so banks should treat provider dependency as a third-party risk, building for portability and keeping sensitive data and model customisation under their own control.

5. TREAT AI RISK AND AI'S FOOTPRINT AS ESG ISSUES – Credibility at Stake

Banks approach responsible AI largely through familiar IT risks, rating data protection and information security highest, while AI-specific risks such as bias and model risk receive less attention. At the same time, banks rate the environmental footprint of AI itself as least relevant, despite rising energy, water and emissions impacts. As AI moves into real decisions and is deployed in the name of sustainability, banks should manage both the AI-specific model risks and the footprint of their own AI use.




IMAGE SOURCE: iStock.com/todamichi, Bild-ID: 2185337747

Introduction

Sustainability factors are increasingly treated as financially material drivers of prudential risk, and supervisors worldwide are tightening their expectations for how banks identify, measure and manage them (UNEP FI, 2026). Therefore, rising regulatory demands, from the CSRD to EBA guidelines and ECB supervisory expectations, require (German) banks to assess and disclose an ever-widening set of sustainability-related risks and opportunities. At the same time, stakeholders such as investors and rating agencies increasingly scrutinise how institutions manage these risks, in order to better assess both their sustainability ambitions and their risk exposure. Not least, banks themselves increasingly want to use ESG information as a steering instrument in its own right, for portfolio decisions, risk management and client engagement, rather than treating it as a compliance exercise alone.

This expanding set of demands comes with rising costs. Regulatory complexity and expanding data volumes, driven by fragmented sources, limited standardisation and rising reporting obligations, place significant strain on data management functions. These constraints are particularly acute in ESG, where much of the relevant information is unstructured, expressed against heterogeneous standards and partially sourced from third parties. Recent simplification efforts, such as the EU Omnibus initiative, have further decreased the expected data availability and quality, especially for small and medium-sized enterprises (SMEs), and therefore added another layer of uncertainty banks must navigate. Data constraints are accordingly cited by 75% of banks as their primary implementation challenge for ESG risk integration, more than double the next most common barrier (UNEP FI, 2026).

Against this backdrop, artificial intelligence has emerged as a key lever. General-purpose AI is already in use across banking functions, so for ESG management the question is no longer whether to use AI, but how to move towards deeper, ESG-specific and responsible integration. AI offers promising avenues to automate data collection and analysis, strengthen risk models, and increase efficiency in reporting, while also opening new perspectives for integrating ESG criteria into credit and investment decisions, helping to transform ESG from a compliance obligation into a strategic instrument.¹

The central question is: Where do banks currently stand in adopting AI for ESG management, what use cases are emerging, and how can this be done responsibly and sustainably?

To this end, this paper draws on focus-group workshops, case studies and a survey conducted among banks participating in the Cluster Initiative, to identify where the sector currently stands, what barriers continue to slow progress and which practical steps can strengthen adoption and responsible implementation. In doing so, the paper is intended as a guide for C-level executives, sustainability, risk and data professionals, and all those working on the integration of AI into ESG management. At the same time, the findings may also inform regulators, policymakers and other stakeholders by highlighting the practical conditions, constraints and governance questions that shape responsible AI adoption in sustainable finance.

¹ Beyond improving ESG management processes, AI may also support sustainability transitions in the real economy. Recent research argues that AI-enabled production systems, predictive maintenance, and intelligent resource management could contribute to more circular business models by improving product longevity, reuse, refurbishment, and resource efficiency. AI may therefore help institutions manage ESG information while simultaneously shaping sustainability outcomes in the sectors they finance (SEB, 2026).

1. AI Adoption Maturity

Context

Before assessing where banks stand, it is useful to clarify what is meant by AI in this paper. AI should not be understood as a single technology, but as a spectrum of applications with different levels of complexity and autonomy. At one end are established tools such as business intelligence and classical machine learning, which support structured analysis and prediction. Further along the spectrum are generative AI systems, which can process language, summarise documents, draft content and support knowledge work. The most advanced current stage is agentic AI, where systems can plan and execute multi-step tasks within defined guardrails. For ESG management, this distinction matters because each step up the spectrum increases both the potential value and the requirements for data quality, expertise, governance and auditability. A chatbot used to summarise sustainability reports raises different questions than an AI system that informs credit decisions or orchestrates ESG workflows across several functions. In this paper, the focus is therefore on generative and emerging agentic AI. Maturity is not judged by the technical label alone, but by what AI is used for, what it affects, and how deeply it is embedded in the organisation and its controls (Fritz-Morgenthal et al., 2022).

The move from isolated pilot projects to organisation-wide implementation faces several practical obstacles. Core banking systems were not built for AI, data sits in silos with no single owner, and the specialist skills to build and validate models are scarce and expensive. Cultural factors compound this. Risk-averse structures slow experimentation, and staff are wary both of relying on AI output and of being displaced by it. These obstacles are not unique to ESG, but ESG makes them more difficult to overcome. Relevant information is often spread across fragmented internal and external data sources, embedded in unstructured documents such as sustainability reports, transition plans and client disclosures, and expressed against heterogeneous standards. At the same time, regulatory requirements continue to evolve, which makes it harder to stabilise processes and scale AI applications with confidence. ESG therefore combines the generic challenges of AI adoption in banking with a particularly demanding data and regulatory environment. The rest of the chapter draws on the workshops as well as a survey of participating banks, to show where they stand.²



IMAGE SOURCE: iStock.com/ismagilov, Bild-ID: 1216960081

² The workshops ran from March to June, while the data for the survey was collected at the end of April 2026.


Implementation Maturity

Findings from the survey describe a sector that has committed to the topic but is still early in operational terms. Within this overall picture, the banks in our survey are spread across the adoption range. Nearly a third (29%) are still gathering information and discussing first use cases, and another 29% are running first pilots. One bank operates individual productive applications in single ESG areas, and a further 29% already run several productive applications that are not yet scaled across functions. Among those not yet at scale, 60% expect to reach it within twelve months, 20% have already started, and 20% have no concrete plans. The commitment is clear, but the operational base today remains limited.³

The clearest pattern concerns the form in which AI is used (**Figure 1**). General-purpose chatbots are in use at every participating bank (100%), whereas

agentic systems have not moved beyond the planning stage. Between them, document-based generative AI and curated prompt libraries are mostly still in development or piloting, and integration into existing workflows is rare. The pattern is consistent. The deeper AI reaches into a process, the lower its maturity, because deeper integration demands more data, more control and more governance than a general chatbot. Beyond the technology, deeper integration also raises a question of responsibility. As AI moves closer to decisions such as credit, who stays accountable for the outcome becomes harder to answer, and such uses face stricter regulatory requirements, which the paper takes up in Chapter 3. Nevertheless, this observation is not unique to ESG. Most banks are still in an early, pilot-driven phase of AI adoption, short of scaled and integrated use (Deloitte, 2025).

Figure 1: Maturity of AI Use in ESG Management, by Form of Use




³ This is broadly consistent with evidence from the wider Swiss sustainable finance market, which points to increasing AI adoption and a growing number of more advanced applications (SSF, 2026a).

Drivers, Barriers and Challenges

The banks also report what drives and what blocks the effective use of AI (**Figure 2**). Only “collaboration between the ESG function, the business and IT” emerges as a real driver (+0.4).⁴ The barriers are more numerous and more pronounced. Technical infrastructure and the existing system landscape stand out by a wide margin (-1.4), followed by limited resources (-0.7) and limited trust in AI output (-0.6). The availability of AI skills is a further,

if milder, barrier (-0.3). Most remaining factors, among them data quality, governance and digital sovereignty, are rather neutral (-0.1). These challenges are not unique to the banks participating in this initiative. A broader survey of German financial institutions similarly identified limited resources, shortages of AI-related skills, regulatory requirements and data-quality issues among the most significant barriers to AI adoption (PwC, 2025).

Figure 2: Drivers and Barriers to Effective AI Use in ESG Management



⁴ The score is the average of all responding institutes' answers on a five-point scale, with each answer weighted from -2 (strong barrier) through 0 (neither barrier nor driver) to +2 (strong driver). A negative value means a factor is, on average, perceived as a barrier and a positive value as a driver, with values near zero indicating a neutral or mixed assessment.

1. AI ADOPTION MATURITY

The close-to-neutral rating of data quality and governance is notable, since both are usually named among the central challenges for AI. A score near zero reflects mixed or muted perceptions rather than low importance. While most banks are still running pilots, data and governance gaps can be worked around, and they become a real constraint only when AI is scaled up. Deloitte (2025) points the same way, calling data the tightest bottleneck for scaling and the governance arrangements of the exploration phase insufficient for it.

The banks do not yet see digital sovereignty as a pressing barrier, but it is nonetheless a key strategic challenge for the sector, and one that external studies already flag for banking more broadly (Deloitte, 2025). One reason for this result may be timing. The survey was conducted in late April 2026 and predates the most recent developments that have since raised the salience of provider dependency.

The concern is concrete. The models and hardware behind current AI come from a small number of mostly US suppliers. A handful of firms provide the leading foundation models, and effectively

one company dominates the specialised hardware they run on. For a bank, this concentration is an operational risk rather than an abstract concern. Cloud-based models raise questions of data outflow and confidentiality. Pricing is likely to rise as the market consolidates, and access itself can become a matter of international politics rather than commerce. This is no longer hypothetical. In mid-2026, access to leading models was restricted on national-security grounds, and some providers temporarily disabled their most capable systems for all customers (Anthropic, 2026). Deploying models locally reduces the data-outflow risk but shifts the dependency onto hardware availability and price, while also creating a risk of falling behind the most recent technological advances. Unfortunately, competitive European alternatives are not yet in sight.

Withdrawing from these providers, however, is not a realistic option. It would leave a bank at a clear competitive disadvantage and, given how heavily AI is now used on both sides of cyber security, at a security disadvantage too. The practical response is not independence but resilience. This is a real dilemma for the sector, and the recommendations below set out how to manage it.




IMAGE SOURCE: iStock.com/Sefa Kart, Bild-ID: 1157017202

Recommendations for Banks

The evidence describes a sector that is committed but uneven, and one that will gain more from AI in ESG by scaling deliberately than by running more pilots. The following priorities emerge from the status quo and the challenges above.

Treat buy-and-build as a portfolio and architecture decision, not a tool-by-tool choice. The survey shows banks weighing each case on its own, buying general-purpose tools for a quick entry and building in-house where they want control, without a clear rule for what should be reused. This matters because case-by-case sourcing risks reproducing the patchwork of disconnected systems banks are still paying to replace in their legacy estate. Banks should therefore buy to enter quickly and build in-house for reusable capabilities, so that each build adds to a shared foundation rather than standing alone.

Tackle the binding constraints, starting with infrastructure and collaboration. Banks identify technical infrastructure and the existing system landscape as the most significant barrier, followed by limited resources. At the same time, they point to close collaboration between ESG, business and IT as the clearest enabler. Banks should build a shared ESG data foundation with clear ownership, connecting the data that today

sits in disconnected silos, which is a core reason infrastructure ranks as their top barrier. This technical foundation should be paired with cross-functional collaboration and involve not only AI and ESG specialists but also the end users and process experts, such as credit analysts, whose work the tools touch. That way capability is spread across teams rather than held in isolated units.

Put provider dependency into the operational-risk framework. The survey does not yet register this as a pressing barrier, but the concentration of models and hardware among a few non-European suppliers is a live operational risk, as the mid-2026 access restrictions showed. Banks should treat it like any other third-party risk. That means building for portability, so that a model or provider can be swapped without re-engineering the surrounding process, keeping sensitive data and any model fine-tuning under their own control, and setting out fallback options before a dependency becomes critical.

Taken together, these steps move a bank from opportunistic experimentation to disciplined scaling. Without them, the sector risks a widening gap between its stated ambition and an operational base that stays shallow, fragmented and dependent.

2. Use Cases

Context

AI use cases in ESG management stretch across the banking value chain, from data analysis and reporting to risk and client-facing tasks. A useful way of structuring these applications is along a “data-to-decision” workflow, spanning data collection, quality assurance, analysis, predictive analytics and content generation (SSF, 2026b). Progress here depends less on the technology than on the discipline to make scarce budgets and developers’ capacity go further. Most use cases rest on a few recurring building blocks, such as data extraction, summarisation, retrieval and document question-and-answer. Banks that build these once and reuse them gain more from each new appli-

cation than banks that commission isolated point solutions (PwC, 2026). Cost efficiency and scale therefore come from reuse, which turns the question of which use cases to pursue into a question of what scales rather than of individual tools. Where that discipline is missing and progress is counted in proofs-of-concept rather than in adoption and measurable impact, initiatives tend to stall at the pilot stage (FIRM, 2025).

What follows sets out how banks decide which use cases to pursue, where they have made most progress, what value they get out of it and concrete examples of use cases.

Selection Criteria, Adoption and Value

When they decide which use cases to pursue, the banks weight technical feasibility most heavily (**Figure 3**).⁵ Reusability of components and scalability lead the ranking, followed closely by internal skills, implementation effort and the expected business case. Fit with existing processes comes next. Strategic relevance and the traceability of a solution rank lower, and data availability and regulatory relevance rank lowest of all. The logic is efficient and consistent with the reuse discipline set out above. The banks favour use cases they can build, connect and scale. At the same time, the results reveal a potential weakness in this approach. The banks place data availability and regulatory fit at the bottom of the list, yet data underpins every use case, and regulatory requirements are the factors most likely to decide whether one lasts and stays compliant.



IMAGE SOURCE: Stock.com/napong rattanakrattiya , Bild-ID: 1929775959


⁵ In Figure 3, the score is the average of all responding institutes’ answers, on a scale from -2 (not relevant) to +2 (highly relevant). The figures are indicative.

2. USE CASES


A further criterion, not explicitly covered in the survey but important for scaling, is user attractiveness. Banks are still at an early stage of their AI journey, and adoption will depend not only on whether a use case is technically feasible, but also on whether employees see it as useful, intuitive and worth integrating into their daily work. Use cases that reduce repetitive work, make expertise easier to access or visibly improve output quality are more likely to build trust and momentum than technically elegant solutions with limited user pull. In that sense, user acceptance should be treated as part of the value case from the start, especially where banks aim to move from pilots to broader adoption.

When it comes to adoption, we find that banks apply AI unevenly across the ESG value chain (**Figure 4**). It is furthest along in analytical tasks, where the work is well defined and the output supports rather than decides. The analysis of external ESG reports and transition plans leads (2.4 on a scale from 0 to 4),⁶ followed by a cluster at 2.0 that covers ESG data collection, knowledge and peer support, ESG-risk documentation, and support for reporting and disclosures. Regulatory analysis, taxonomy and activity mapping, external monitoring and ESG risk in the credit process sit in the middle (1.4 to 1.6), mostly in development or piloting. The fields that reach deepest into decisions and clients lag furthest behind.

Figure 3: Relevance of Criteria for Prioritising AI Use Cases



⁶ In Figures 4 and 5, the score is the average of all responding institutes' answers on a scale from 0 to 4. For the stage of implementation (Figure 4), 0 means no use case is in place and 4 means the field is rolled out across the bank. For realised value (Figure 5), 0 means no value and 4 means very high value. The figures are indicative.

Figure 4: Current Implementation Status of AI across ESG Areas of Use

Portfolio analysis (1.0), the integration of ESG factors into decisions and processes (0.7), and client engagement and advisory (0.3) are, in most banks, either at an early stage or not yet on the agenda. The pattern mirrors the maturity finding from the previous chapter. The more a use case moves from supporting analysis to shaping a decision, the less advanced it is.

The findings on client engagement and advisory (and related value creation, which we discuss in

the next paragraph) should be interpreted with caution. There are indications that AI use cases are already being developed across financial services for targeted communication, onboarding, intelligent servicing, advisor support and personalised client orchestration (Capgemini, 2025a; Capgemini, 2026). The low maturity observed here may therefore reflect less a lack of client-facing AI adoption of use cases overall (which we did not assess) than the fact that these are probably developed without a specific ESG lens. The opportunity

2. USE CASES

for ESG management is to pilot these broader client-facing AI patterns with an ESG-specific lens, for example in transition-plan dialogue, ESG product advice or sustainability-related client engagement.

The value the banks report from this is real but concentrated (**Figure 5**). It shows up most in learning to work with AI (2.4), in time saved (2.3), in more consistent and comparable analysis (2.0), and in better quality of output (1.9). It shows up least in the areas that would generate new business, such as enabling new services (0.3), reaching new client segments (0.4) or developing new products (0.7), where almost no bank reports meaningful value yet. The gains, in other words, are in efficiency and capability rather than in revenue.

How that value is tracked is the weaker point. More than half the banks (57%) do not yet capture the value of their ESG-AI systematically, and none rely mainly on quantitative indicators. This makes it hard to compare use cases and to know which pilots are worth scaling. At least one bank sets a higher bar, requiring the expected value of a tool to be shown and quantified before development and running the same testing process it applies to any other application before roll-out. Making that the norm rather than the exception would give the selection and scaling decisions in the rest of this chapter a firmer basis.


The survey does not record why value measurement remains limited, but workshop discussions and broader evidence on AI adoption in business operations suggest a possible explanation. Banks may currently prioritise rapid adoption and practical use of AI over systematic measurement of efficiency gains, partly because formal measurement can make the organisational implications of AI more visible. If time savings are quantified too explicitly, employees may read this less as evidence of process improvement and more as a signal of future role reduction, reinforcing existing concerns about displacement. In that sense, premature



IMAGE SOURCE: iStock.com/metamorworks, Bild-ID: 1321462048

or poorly communicated value measurement could create resistance and slow implementation (Capgemini, 2025b).

The challenge is therefore to measure value in a way that supports adoption rather than undermines it. More systematic value tracking should frame AI benefits not only as cost reductions, but also as better work, higher quality, reduced manual burden and redeployed capacity. The metrics should then cover both generic AI benefits and ESG-specific outcomes. Efficiency indicators could include processing time, manual effort saved, cost per analysis or reporting cycle time. Quality indicators could cover error reduction, consistency of outputs, completeness of data extraction, traceability of sources and the share of outputs accepted without major rework. For ESG use cases, the value case should also capture whether AI improves the substance of ESG management, for example by identifying risks earlier, improving portfolio transparency, supporting transition analysis or strengthening client dialogue. A disciplined approach would define these indicators before development, test them before roll-out and monitor them continuously once the tool is in use. This would turn value tracking from an ex post justification into part of the use-case design itself.

Figure 5: Value Realised from AI in ESG Management

Overview of use cases and case studies

The previous section showed that banks are beginning to apply AI across a broad range of ESG activities, but that adoption remains uneven and value is still concentrated in analytical and efficiency-oriented tasks. The use cases below make this pattern more concrete. They show where AI is already being tested or used in practice, how recurring AI patterns such as retrieval, extraction, classification, monitoring and source-backed text generation appear across different banking processes, and where the main implementation and governance questions arise.

Table 1 provides a comparative overview of four selected use cases presented in the Cluster workshops. It is not intended as a ranking or as a complete market inventory. Instead, it structures the cases along the ESG activity addressed, the underlying pain point, the AI pattern used, the data and workflow involved, the value created, and the governance conditions that shape responsible use. For three of the four use cases, the overview

is complemented by case studies that go one level deeper, describing the practical starting point, AI-supported workflow, human oversight, benefits, limitations and scaling potential in more detail. A broader catalogue of additional and potential AI use cases for ESG management is provided in the Appendix.

Across the use cases and case studies, several common conclusions emerge. AI is most advanced where it supports data, document and analytical work, helping experts reduce manual effort, improve consistency and navigate heterogeneous information without taking over final decisions. The case studies make this point more tangible. Whether in ESG credit analysis, reputational risk monitoring or greenwashing screening, AI creates value by narrowing the field of attention, structuring evidence and making source material easier to use, while the final judgement remains with experts. The same AI patterns also recur across different ESG topics. Retrieval, extraction, classification, monitoring and source-backed text generation appear in credit analysis, portfolio transparency, reputational risk monitoring and greenwashing screening alike, reinforcing the case for reusable capabilities rather than isolated tools.

At the same time, the cases show that AI creates durable value only when embedded in controlled workflows, with human review, source traceability, clear classification logic and documented limitations. Scaling therefore depends less on model sophistication than on data access, process integration, user trust and governance. The cases confirm that AI in ESG management is already practically useful, but its durable value will come from disciplined implementation rather than experimentation alone.



IMAGE SOURCE: iStock.com/WANAN YOSSINGKUM, Bild-ID: 1922342677

2. USE CASES

Table 1: Overview of AI Use Cases in ESG Management

ESG activity and banking process	Pain point	AI pattern	Data sources and workflow	Primary value	Governance and limitations	Current maturity / next step
Helaba & Deka Bank: AI-supported ESG analysis in the credit process						
ESG risk assessment in credit analysis, including support for ESG risk scores and transition assessments.	Analysts face high manual effort, inconsistent information formats and rising requirements for more detailed ESG risk assessments.	Grounded information retrieval and source-backed text generation.	The analyst creates a document-based knowledge base, selects the relevant assessment type, receives AI-supported answers and then checks the results for plausibility.	Time savings, reduced analyst workload, better source traceability and more consistent preparation of ESG assessments.	The tool is positioned as analyst support rather than replacement. Outputs are checked by humans and linked to sources and page references.	Productive no-brainer case; next step is further refinement and scaling through application to additional analyses, even beyond ESG matters
Novata: AI-supported reputational risk assessment						
Monitoring of ESG controversies and reputational risks for portfolio companies, counterparties and supply chains.	Manual monitoring does not scale across many counterparties, while static tools can miss fast-developing controversies or generate excessive noise.	Web search, multilingual retrieval, filtering, risk scoring, alerting and dashboard-based monitoring.	The system scans media, regulatory filings and enforcement signals, filters them according to configured risk categories and presents prioritised risk signals in a dashboard.	Continuous monitoring, earlier detection of controversies, prioritised analyst attention and a more structured view of reputational risk exposure.	Current limitations include web-search coverage, especially for private companies, and efficiency constraints when retrieval and assessment are combined in one step.	Existing solution with further development towards specialised pipelines, structured event feeds and evaluation of alternative large language models.

2. USE CASES

ESG activity and banking process	Pain point	AI pattern	Data sources and workflow	Primary value	Governance and limitations	Current maturity / next step
----------------------------------	------------	------------	---------------------------	---------------	----------------------------	------------------------------

Commerzbank: AI-supported screening of potential greenwashing risks under EmpCo requirements

Review of public web content for environmental and sustainability-related claims linked to EmpCo requirements and their implementation into the German UWG	Large volumes of web content make comprehensive and recurring manual reviews difficult, particularly where relevant claims are spread across many web-pages.	Web content retrieval, rule-based term screening and generative AI-assisted contextual assessment.	Public webpages are collected, screened for predefined environmental and sustainability-related terms, and assessed against EmpCo-related dimensions to prioritise pages for expert review.	Enables risk-based prioritisation, improves transparency across large volumes of content and helps identify potentially relevant claims that may be difficult to find manually.	The tool provides indicative screening results only. Expert review remains necessary, and criteria may need to be updated as regulatory interpretation evolves.	Implemented as a screening and prioritisation tool; potential next steps include extending coverage to additional communication channels and embedding it in recurring review processes.
--	--	--	---	---	---	--

DZ BANK: Activity splitting and SDG classification

Portfolio transparency, SDG classification and broader sustainability metrics in the corporate credit portfolio.	The creation and updating of activity splits is highly manual, relies on heterogeneous public information and is difficult to scale across large portfolios.	Information retrieval, activity mapping, classification and potentially agentic process automation.	Public sources such as annual reports, sustainability reports, company websites and public registers are used to derive revenue-based activity profiles, which inform SDG classification and other sustainability metrics.	Better portfolio transparency, improved data quality for downstream methods, more scalable classification and synergies for climate alignment, CO ² accounting and reporting.	Public information can be incomplete, unclear or outdated. Human review, four-eyes controls and clear classification logic remain important.	Manual process with clear automation potential; a digitalisation project is planned as a possible path towards more automated operation.
--	--	---	--	--	--	--



HELABA CASE STUDY



AI-supported ESG analysis in the credit process

Extracting ESG information from large documents to support ESG risk scores and transition assessments

Starting point and challenge

DokExtrakt is an internal AI application that supports ESG analyses by extracting relevant information from extensive documents such as sustainability reports and annual financial statements. It is designed to facilitate key components of the ESG credit process, including the ESG risk score and transition assessment frameworks. Users upload ESG-related documents and receive structured outputs that can be integrated directly into their analytical workflows. The main users are credit analysts and staff involved in credit processes who need to systematically identify and process ESG information from large document sets.

Before DokExtrakt, ESG information had to be identified manually in sustainability reports, annual financial statements and other ESG documents. Relevant passages needed to be found, interpreted and then transferred into internal templates. This manual approach is particularly time-consuming and error-prone when dealing with documents that can reach several hundred pages – and sustainability reports have grown significantly in length in recent years. In addition, analysts typically cover different sectors, and although continuous training is provided on ESG assessments, it remains challenging to perform this task consistently. The need to process large volumes of ESG documentation efficiently, while maintaining the quality and consistency of ESG data for credit processes, created demand for an AI-based support.

AI-supported workflow and human oversight

Extracting information from documents is a common feature of many AI applications, and Helaba already operates its own AI platform. So what makes DokExtrakt special? For every single piece of information required to conduct an ESG risk assessment, ESG and AI specialists have worked together to design high-performing prompts that provide contextual information and clear guidance on the desired output. These prompts run in the background, so users do not have to think about prompt design at all.

The workflow has been greatly simplified and now consists of just four clicks:

1. Upload documents as a knowledge base from which the AI can extract information.
2. Select the type of ESG risk assessment to be conducted.
3. Start the analysis.
4. Download the outputs.

► Continue reading on the next page.



Of course, the work does not end there, as plausibility checks remain a crucial step in the process. An important aspect is that the user – typically the credit analyst – remains responsible for the content when it is used in the credit process. This ensures that AI-generated content is subject to expert judgement before it integrates into risk assessments or credit decisions. Designing the prompts was an iterative process and the result of close collaboration between ESG, AI and credit experts. This interdisciplinary approach was one of the key success factors for implementation.

Benefits, limitations and scaling potential

The application has been very well received by users, particularly because of the time saved per ESG assessment. As the workflow was designed so that no AI expertise is required, user acceptance was high from the outset.

While DokExtrakt was initially developed for ESG risk assessments, we quickly realised that its capabilities are relevant for many additional use cases. Scaling the application is therefore an important ongoing objective. We aim to continuously improve the prompts in the background in line with evolving requirements from our ESG risk assessment frameworks and sustainability-related disclosures.



NOVATA CASE STUDY



AI-supported reputational risk assessment

Monitoring ESG events across portfolio companies, counterparties and supply chains

Starting point and challenge

A single reputational controversy can change how a financial institution views a company. A supplier being named in a forced-labour investigation, a borrower who is facing environmental fines, a counterparty whose founder is under criminal inquiry: each is a reputational event, and each can be an early sign of a company's credit, ESG or compliance risk well before it reaches a financial statement. The challenge of monitoring reputational events that lead to serious ESG or compliance risks is that the information is scattered across thousands of sources, spans many languages, and is difficult to track manually across a large book of names.

AI-supported workflow and human oversight

Novata, a sustainability data and technology provider, developed the reputational-risk module within its Risk Atlas product to address this. It is not a passive news feed but rather an actionable tool that enables users to monitor a defined set of companies, whether portfolio holdings, counterparties, or suppliers, and to separate the events that matter from the noise around them.

The approach separates discovery from judgement, and the AI-supported discovery process is built for scale. A recall-first search layer continuously scans news media, regulatory filings and enforcement records, including the national and local-language sources that generic feeds tend to miss, and groups related coverage into distinct events rather than repeated headlines. For companies with little public reporting, a complementary search step looks specifically for what the mainstream sources overlook, historically the weakest point for privately held names.

Next, an internal materiality model reviews each event: how credible the source is, whether the company is actually the subject, whether the event is new, and whether it points to a financial consequence, then labels it by materiality with a confidence level. Taking a phased approach is important in order to deliver a robust solution: discovery has to be exhaustive, since missing a material event is costly; assessment has to be selective, since no analyst can act on everything. Asking a single AI action to accomplish both tends to trade recall against precision, and obscures why any item surfaced. It is also why large-language models are used only to summarise and interpret the results, never to find the events in the first place.

► Continue reading on the next page.



Each event is then mapped to the frameworks institutions already work with, principally the UN Global Compact and the OECD Guidelines. This is what makes the output usable rather than merely interesting. A raw negative-news flag does not tell a risk officer what to do; a labour-rights allegation tied to a UN Global Compact principle, or a due-diligence failure framed against the OECD Guidelines, lines up directly with the language of a bank's own exclusion policy, its ESG credit assessment, or a sustainability-linked loan condition. It also lets an institution filter by the norms it has committed to and report in those same terms.

Throughout the process, Risk Atlas' reputational risk monitoring serves as analyst support, not an automated verdict. Every event carries its original source, so a reviewer can trace a finding back, judge the outlet and set aside false positives before anything is acted on. A financial institution can focus on a single, continuously updated and prioritised view of reputational exposure across its portfolio, with company and portfolio-level roll-ups.

Benefits, limitations and scaling potential

The module is designed to surface risks that have been publicly reported, which means it cannot identify events that have not yet entered the public domain. If, for example, a company is experiencing layoffs, a child labor issue, or a cyberattack that has not been publicly reported, the module will not pick it up. Its outputs are therefore indicative rather than a definitive risk score, and the capability is being developed iteratively rather than presented as a finished, real-time platform. Its central focus is addressing a problem the market has long struggled with: surfacing credible, framework-aligned signals for smaller and private companies where public information is often thin.

Although this case is framed for banking, the same module is used by Novata's investor clients to monitor their own portfolios, and because it sits within Novata's wider ESG platform, the risk results can feed back into the sustainability work already underway there (for example in deciding which ESG data points are worth collecting from a company in the next reporting cycle).



COMMERZBANK CASE STUDY



AI-supported screening of potential greenwashing risks under EmpCo requirements

Identifying and prioritising environmental claims in public web content for expert review

Starting point and challenge

This case study describes how Commerzbank uses AI-supported screening to identify potentially critical environmental and sustainability-related claims in public web content. The use case was triggered by the Empowering Consumers for the Green Transition (EmpCo) Directive and its implementation into the German Act Against Unfair Competition (Gesetz gegen den unlauteren Wettbewerb, UWG). These requirements increased the need to review public-facing environmental claims in a structured and repeatable way. The solution is designed to support Sustainability Management, legal experts and business units responsible for external web content by creating a risk-based starting point for expert review rather than an automated compliance decision.

Before the tool was introduced, potentially relevant claims had to be identified manually across many webpages. This made comprehensive and recurring reviews resource-intensive, particularly where claims were spread across different areas of the website or embedded in less visible sections. At the same time, a purely manual approach made it difficult to ensure consistent coverage across large volumes of content. The main challenge was therefore not only to find individual keywords, but to help experts understand where potentially relevant statements appeared, how they related to EmpCo-related dimensions and which pages should be prioritized in the review.

AI-supported workflow and human oversight

The workflow combines web content retrieval, rule-based term screening and generative AI-assisted contextual assessment. Public webpages are collected and screened for predefined environmental and sustainability-related terms, such as references to climate neutrality, environmental benefits, sustainability commitments or environmental labels. Generative AI then assesses the context of the identified content against predefined EmpCo-related assessment dimensions. Rather than treating every keyword hit as equally relevant, the model supports the classification and prioritisation of webpages according to their potential relevance for further review.

► Continue reading on the next page.



COMMERZBANK

The output is a structured overview of assessed webpages. This helps reviewers move from broad website coverage to targeted expert assessment and makes the review process more transparent. Importantly, the tool separates initial screening from final and comprehensive judgement: it helps identify and prioritise potentially relevant content, while the interpretation of individual claims remains with the responsible experts.

Benefits, limitations and scaling potential

In practice, the main value lies in improving transparency across large volumes of public web content and helping reviewers allocate their attention more efficiently. The tool can surface potentially relevant environmental claims that may be difficult to locate through manual screening alone, while also creating a more consistent basis for deciding which pages should be reviewed first. This supports a more efficient use of expert capacity without removing the need for comprehensive judgement.

The limitations of the approach are closely linked to its role as a prioritisation instrument. The results are indicative and depend on the quality of the underlying website collection, the predefined screening terms and the assessment criteria applied. False positives and false negatives may occur, and environmental claims often require context-specific interpretation. In addition, regulatory expectations may evolve as market practice, guidance and case law develop. The screening logic and assessment dimensions therefore need to be kept under review rather than treated as fixed parameters.

Potential next steps include extending the approach to additional communication channels and content formats and embedding it in recurring review processes. More broadly, the same pattern of retrieval, filtering, AI-assisted contextual assessment and prioritisation could be applied to other sustainability or regulatory use cases that require the review of large volumes of unstructured information. For broader scaling, the key condition is to maintain clear responsibilities: AI can support coverage, consistency and prioritisation, but accountability for final interpretation and follow-up action should remain with the relevant expert functions.

Recommendations for Banks

The evidence describes a sector that chooses use cases largely for what is quickest to build, while the case studies show that durable value depends on whether AI can be embedded in controlled workflows, trusted by users and measured in practice. The following priorities emerge from the findings above.

Widen the criteria for choosing use cases beyond technical feasibility. The banks select mainly for technical feasibility and reuse, while ranking data availability, regulatory fit and strategic relevance lowest. The case studies show why this is too narrow. Use cases last when they combine a clear pain point with accessible source material, defined workflow integration, human review, traceability and a realistic path to scaling. Banks should therefore give these implementation conditions real weight from the start, so that effort goes not only to what is quick to build, but to what can be governed, adopted and reused over time.

Prove value before scaling, and measure it in terms users recognise. Banks should require a clear value case before building and test it at roll-out, using both quantitative and qualitative indicators. The case studies make the value case more concrete. AI can save time in ESG credit analysis, prioritise analyst attention in reputational risk monitoring, improve transparency across large volumes of web content and create more consistent, source-backed outputs. Value tracking should therefore not be framed only as cost reduction. It should also capture higher output quality, reduced manual

burden, redeployed capacity, better traceability and ESG-specific outcomes, such as earlier risk identification, improved portfolio transparency or stronger preparation for client dialogue. This makes scaling decisions more evidence-based without turning measurement into a source of employee resistance.

Match ambition to readiness, and keep expert judgement clearly separate from AI support.

Banks should scale analytical and data-near use cases first, especially where AI can narrow the field of attention, extract or structure evidence and prepare outputs without taking over final decisions. The case studies show that this separation is a core design principle: AI supports screening, retrieval, prioritisation and drafting, while experts remain responsible for interpretation and follow-up. Decision-critical and more agentic applications should therefore be pursued only where ESG data quality, source traceability, controls and accountability are sufficiently mature. Client-facing ESG applications should not be treated as out of scope, but tested selectively where existing advisory or engagement AI use cases can be extended with an ESG-specific lens. Pushing depth ahead of readiness risks stranded pilots, while ignoring client-facing experimentation risks keeping AI value confined to efficiency gains.

Taken together, these steps turn the chapter's question from what can be built into what is worth building and able to last. Without them, the sector keeps generating pilots faster than it can prove or scale them.

3. Responsible & Sustainable AI Governance

Context

AI in ESG management is no longer confined to back-office support. It is now entering regulated, high-stakes processes such as credit decisions, disclosure and reporting, where a wrong or unexplained output has real consequences. This places governance at the centre. The outputs must be trustworthy, and a person or function must remain accountable for them (Fritz-Morgenthal et al., 2022). Recent assessments likewise note that AI capabilities are advancing faster than existing approaches to measurement, evaluation and governance, reinforcing the need for robust arrangements (United Nations, 2026). The regulatory framework is already taking shape. Under the EU AI Act, uses such as credit scoring count as high-risk and carry obligations for transparency, human oversight and documentation (European Union, 2024).

Responsible use is only one lens. AI also carries an environmental cost of its own, which is the chapter's second concern. The compute behind modern AI consumes energy and water and relies on hardware whose production and the build-out of data centres generate emissions. Research highlights the broader environmental footprint of AI, including embodied emissions within semiconductor supply chains, electronic waste and resource extraction associated with hardware production (Oxford Sustainable Finance Group & Theia Finance Labs, 2026). Yet the sustainability implications of AI are not limited to its own footprint. For banks, this footprint arises through two channels. One is the bank's own use of AI. The other, still emerging as a portfolio question, is the foot-

print of the AI and data-centre infrastructure that banks finance. This chapter concentrates on the bank's own use.⁷

What follows sets out how far the banks have taken responsible-AI practices, what the evidence says about the environmental footprint of AI, and what both mean for how they should proceed.



IMAGE SOURCE: iStock.com/Miradeo, Bild-ID: 58529860

⁷ The financing channel is set aside here because it is a portfolio and credit-risk question that goes beyond this paper's focus on the bank's own use of AI.

Responsible AI


Asked which aspects of responsible AI matter to them in principle, the banks put data protection and information security first (**Figure 6**), followed closely by explainability, transparency and auditability, and by human oversight of AI output.⁸ Documentation of sources and decisions, the management of bias and model risk, and clear responsibilities and governance come next, with stakeholder involvement lower down. The ranking is worth a second look. Banks lead with data protection and information security, which they know well from general IT, while the risks that are new with AI, bias and model risk, sit lower, and no bank rated them highly relevant. One aspect sits well below the rest, and the paper returns to it under sustainability.

Responsibility for AI in ESG is not yet settled. A hybrid model that combines the ESG function

with central or other units is the most common arrangement (43%), while the remaining banks are split evenly across centralised, decentralised and dedicated-AI-office setups, and some have no clearly defined ownership at all. No single standard has emerged. External guidance and industry experience converge on a hybrid governance model, anchored in a central AI office or centre of practice and integrated into existing organisational and regulatory structures rather than managed through parallel frameworks (Deloitte, 2025; PwC, 2025).

The practices the banks discuss reflect a specific property of the current technology. Large language models are probabilistic rather than deterministic, so the same prompt can produce different results (FIRM, 2025). The mitigations raised are practical. Accuracy-critical outputs are checked

Figure 6: Relevance of Responsible-AI Aspects



⁸ The score is the average of all responding institutes' answers on a scale from -2 (not relevant) to +2 (highly relevant). The figures are indicative.

against the original sources, the number of automated AI steps is kept low so that a person stays in the loop, and prompts, models and results are documented so that a decision can be reconstructed. These echo the transparency and oversight aspects the banks rate highly and the documentation duties the EU AI Act imposes (European Union, 2024).

The management of bias and model risk, which the ranking names but the section has not yet examined, is specific to how AI works. Unlike conventional software, AI models learn from data, so they can carry forward bias present in that data, and their behaviour can shift as data and use change. This makes AI a matter of model risk as much as of IT security, and the controls banks already apply to other models (e.g., validation before deployment, monitoring once in use, and clear ownership) extend to it. These matter more as AI moves from supporting analysis towards shaping decisions such as credit, where a biased or unstable model has real consequences. This becomes even more important for agentic systems, where autonomy must be bounded by clear scope, traceability, explainability and human intervention rights (Capgemini, 2025c).

Sustainability of AI

The increasing use of AI will undoubtedly have significant environmental impacts. Electricity demand from data centres is set to roughly double by 2030, and AI already accounts for an estimated 15% to 20% of it, a share that could approach 40% by the end of the decade (Allianz, 2026). Most of the burden comes not from training models once but from running them. Inference, the continuous use of AI to answer everyday requests, is estimated to make up 80% to 90% of AI's energy footprint (UNU-IN-WEH, 2026). And the cost extends beyond electricity. Data centres consumed some 814 billion litres of water in 2025 and could need 1.3 to 1.8 trillion litres by 2030, comparable to Switzerland's annual

use. Their carbon footprint of around 286 million tonnes of CO₂ equivalent in 2025 could more than double by 2030 if the grid is not cleaned in step (Allianz, 2026). This footprint is driven by AI use across the economy, not by banks alone. Yet for a bank that deploys AI in general but also in the name of sustainability, its own use is still potentially a relevant issue for its own ESG management.

Strikingly, though, the banks place the environmental footprint of AI itself at the very bottom of the responsible-AI list. It is the only aspect they lean towards rating as not relevant (-0.6), which suggests the topic is barely on their radar while attention stays on rolling AI out rather than on managing its impact. The survey does not explain why the banks rank the impact so low, but several explanations are plausible. Keeping pace with the technology, given that not using AI is not a realistic option, can push its own footprint down the agenda. This may change once AI-related computing, cloud and infrastructure services become more visible in sustainability reporting, for example as part of Scope 3 accounting for purchased goods and services. At that point, the footprint of AI would no longer be only an internal efficiency or IT question, but part of the bank's reported sustainability performance. Yet, a major barrier to managing the footprint of AI usage is that there is so far no established way to measure it. Promisingly, though, international efforts are underway to develop harmonised indicators and methodologies for assessing the environmental impacts of AI systems (Coalition for Sustainable AI, 2026).

Nevertheless, where it is not measured directly, cost serves as a rough proxy, since larger, newer models and more tokens tend to use more energy. Measures introduced for cost efficiency, such as flagging model costs to steer users towards lighter ones for simple queries, can therefore reduce it as well. Interestingly, experimental evidence suggests that substantial reductions are achievable without compromising performance. Techniques such as model quantisation, prompt optimisation and the

3. RESPONSIBLE & SUSTAINABLE AI GOVERNANCE

use of smaller task-specific models reduced energy consumption by up to 44%, 50% and 90% respectively in controlled testing (UNESCO, 2025). This is a promising lever that users of AI, including banks, should be aware of, also since costs for using AI have been increasing more and more recently.

At the same time, the picture is not one-directional.⁹ AI is increasingly being applied to improve the efficiency of energy, water and other resource-intensive systems and could, if deployed effectively across the economy, reduce emissions by more than its own infrastructure adds. This potential, however, does not remove the need to manage AI's own footprint. Assessments of sustainable AI should therefore consider both sides: the contribution AI can make to sustainability outcomes and the rising environmental costs of running and building AI systems, including the embodied emissions in servers and chips (Salesforce, 2025; Allianz, 2026).

For banks, this means that the sustainability assessment of AI should not stop at model choice or computing intensity. It also depends on the infrastructure behind the service. Relevant questions include whether the bank or its providers use low-carbon electricity, whether data centres are designed for circularity and resource efficiency, and whether waste heat and cooling water can be reused rather than treated only as environmental burdens. In cities such as Frankfurt, the reuse of data-centre waste heat is increasingly discussed as part of municipal heat planning, while the decarbonisation of electricity and heating systems can reduce the footprint of AI infrastructure over time. These developments do not eliminate the environmental cost of AI, but they show that the footprint is shaped by energy procurement, local infrastructure and circular design choices as much as by the AI model itself.



IMAGE SOURCE: dzmitryk87_316679226 - stock.adobe.com

⁹ See this link for an inventory of reports on Sustainable AI, covering both of those aspects: <https://www.sustainableaicoalition.org/inventory-of-key-reports-on-sustainable-ai/>

Recommendations for Banks

The evidence describes responsible-AI intent that governance has not yet caught up with, and an environmental cost the banks barely track. The following priorities emerge from the findings above.

Give AI in ESG a clear owner and a governance frame aligned to the EU AI Act. Responsibility is still split across hybrid, central and decentralised setups, and some banks have no defined ownership at all. Banks should settle this now, anchoring a hybrid model in a central AI office or centre of practice with defined roles and clear ownership of model risk. They need to then align it to the EU AI Act, inventorying high-risk uses such as credit scoring, documenting them and ensuring human oversight.

Build responsible-AI practice into the workflow, not around it. The banks already name the right safeguards. They check outputs against sources, keep a person in the loop, and document prompts, models and results. But they treat these as ad-hoc fixes for the fact that large language models can return different results from the same prompt. Banks should make them standard for accuracy-critical ESG uses, so that every material output stays reconstructable and auditable.

Manage AI's own footprint rather than ignore it.

The banks rate the environmental cost of AI the least relevant responsible-AI aspect, yet the evidence shows a footprint that is large and rising. A bank that deploys AI in the name of sustainability should measure and manage its own AI footprint, drawing on emerging international standardisation efforts as they mature. That means right-sizing models to the task rather than defaulting to the heaviest, and favouring more efficient, lower-carbon infrastructure to run them on.

Taken together, these steps make responsible and sustainable AI something the banks manage rather than merely endorse. Without them, the gap between what banks say about AI and how they run it will keep widening.

Outlook: From Scaling to Agentic AI

This paper has focused on how banks move from pilots to scaling. What comes after scaling is agentic AI, systems that plan and carry out multi-step tasks on their own within set guardrails. It is the least mature form of AI in ESG and the hardest to audit, so its adoption will hinge on the governance questions raised in **Chapter 3**. For now it barely exists in production, still mostly planned rather than live, and the first concrete moves are isolated, such as turning an existing use case into an agentic one. Outside ESG, current examples of agentic AI in financial services are concentrated in risk and compliance processes such as client onboarding, credit analysis, AML investigations and negative-news screening (UiPath, 2025).

In the near term, the work is less about agents than about discipline. It means scaling the use cases that already deliver, integrating them so they stay auditable and within the rules, and building AI and ESG skills across the bank. AI agents for day-to-day ESG are a stated aim, and one proposed route combines robotic process automation with agentic AI, an approach some call agentic process automation. Broader research on AI in business operations similarly stresses that agentic AI delivers value only when embedded in redesigned processes and governed workflows, rather than added to isolated tasks (Capgemini, 2025b).

Two dilemmas will shape this next phase. One is external. Reliance on a few non-European providers of models and hardware remains an unresolved strategic and operational risk (**Chapter 1**). Withdrawing is not realistic, given the competitive and security cost, so the task is to manage the dependency rather than escape it, through portability and by keeping sensitive data and model customisation under the bank's own control. The other is internal. As systems act more autonomously, they can find loopholes and work around the controls meant to contain them, so more autonomy brings both more value and more risk (**Chapter 3**; Capgemini, 2025c). Here too the task is not to avoid autonomy but to keep agentic systems within their guardrails, which is as much a risk-management question as a technical one. Autonomy also cuts the other way. Agentic systems can expose inefficiencies and reshape roles, a workforce and change-management question banks will need to handle as carefully as the technical controls. In both cases, the lesson is the one that runs through this paper. The gains from AI in ESG will depend not on how advanced the technology becomes, but on how deliberately banks scale it and how well they keep it in check.



IMAGE SOURCE: iStock.com/Winxclub, Bild-ID: 219141820

References

Allianz Trade (2026)

Code, carbon, kilowatts:
AI's hidden toll and the race to green the grid
Allianz Research, 30 June 2026
[2026-06-30-ai-code-carbon-AZ.pdf](#)

Anthropic (2026)

Statement on the US government directive
to suspend access to Fable 5 and Mythos 5
<https://www.anthropic.com/news/fable-mythos-access>

Capgemini (2025a)

Attract, engage, and delight:
Spin the customer centricity flywheel
World Report Series 2025: Retail Banking
https://www.capgemini.com/wp-content/uploads/2025/03/WRBR_Final_25_webblock-version-1.pdf

Capgemini (2025b)

AI in action:
How gen AI and agentic AI redefine business operations
https://www.capgemini.com/wp-content/uploads/2025/07/30062025-Digital-AI-in-Business-Operation-CRI_V6.pdf

Capgemini (2025c)

Rise of agentic AI:
How trust is the key to human-AI collaboration
https://www.capgemini.com/wp-content/uploads/2026/02/AI-Agents_web_160226-1.pdf

Capgemini (2026)

Wealth.AI:
Unlocking personalization through
augmented intelligence
World Report Series 2026: Wealth Management
<https://www.capgemini.com/insights/research-library/world-wealth-report/>

Coalition for Sustainable AI (2026)

Standardization for AI Environmental Sustainability:
Towards a coordinated global approach
Updated version published for the
AI Impact Summit 2026
https://www.sustainableaicoalition.org/wp-content/uploads/Global_Approach_AI_Summit_2026-Update_FINAL.pdf

Deloitte (2025)

Künstliche Intelligenz im Bankensektor:
Wie Banken ihre Potenziale aktivieren
und Profitabilität steigern
Deloitte
[Künstliche Intelligenz im Bankensektor | Deloitte Deutschland](#)

European Union (2024)

Regulation (EU) 2024/1689 of the European Parliament
and of the Council of 13 June 2024 laying down
harmonised rules on artificial intelligence
(Artificial Intelligence Act)
Official Journal of the European Union, L 2024/1689,
12 July 2024
<http://data.europa.eu/eli/reg/2024/1689/oj>

Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022)

Financial Risk Management and Explainable,
Trustworthy, Responsible AI
Frontiers in Artificial Intelligence, 5, Article 779799
<https://doi.org/10.3389/frai.2022.779799>

FIRM – Frankfurt Institute for Risk Management and Regulation (2025)

Agentic AI in Risk & Compliance:
Von Automatisierung zu intelligenter Governance
White Paper, Round Table Artificial Intelligence
(authors: S. Fritz-Morgenthal, J. Papenbrock,
M. Währisch)
https://firm.fm/wp-content/uploads/2025/11/FIRM_WP_Agentic_AI_final.pdf

Oxford Sustainable Finance Group & Theia Finance Labs (2026)

Silicon Boundaries: Understanding, anticipating,
and managing the societal instability that can be caused
by AI and related technologies
<https://theiafinance.org/wp-content/uploads/2026/03/Silicon-Boundaries-Deck-v6-lo-res.pdf>

PwC (2025)

Einblicke zur Künstlichen Intelligenz im deutschen
Finanzsektor: PwC's deutschlandweite KI-Marktstudie
im Finanzsektor
<https://www.pwc.de/de/finanzdienstleistungen/pwc-studie-ki-im-finanzsektor-2025.pdf>

Salesforce (2025)

AI Sustainability Outlook:
The Challenges, Potential, and Path Forward
<https://www.salesforce.com/en-us/wp-content/uploads/sites/4/documents/company/sustainability/salesforce-ai-sustainability-outlook.pdf>

SEB (2026)

Sustainable Finance Outlook:
Making Circular Economy Bankable
<https://sebgroup.com/siteassets/documents/press/cision/documents/2026/20260218-sebs-sustainable-finance-outlook-making-circular-economy-bankable-en-gb-0-3719284-5307994.pdf>

SSF (2026a)

Swiss Sustainable Investment Market Study 2026
<https://www.sustainablefinance.ch/api/rm/35M4WBJ83Q97M35/ssf-2026-investment-market-study-final.pdf>

SSF (2026b)

How Artificial Intelligence Contributes to
Effective Sustainable Finance
<https://www.sustainablefinance.ch/api/rm/V56B4QV52558984/20260129-ssf-ai-and-sustainable-finance.pdf>

UiPath (2025)

Real-world agentic AI use cases in financial services:
Cutting through the noise with actual client examples
<https://www.uipath.com/resources/automation-whitepapers/real-world-agentic-ai-use-cases-in-financial-services>

UNEP FI (2026)

The Landscape of Sustainability Risk Integration:
The current state of sustainability risk management
practices in banks
<https://www.unepfi.org/industries/banking/the-landscape-of-sustainability-risk-integration/>

UNESCO (2025)

Smarter, Smaller, Stronger:
Resource-Efficient Generative AI & the Future
of Digital Transformation
<https://unesdoc.unesco.org/ark:/48223/pf0000394521>

United Nations (2026)

Preliminary Report of the Independent
International Scientific Panel on AI:
Evidence-based assessment of opportunities,
risks and impacts of artificial intelligence
https://www.un.org/independent-international-scientific-panel-ai/sites/default/files/2026-07/en_Preliminary%20Report_.pdf

UNU-INWEH (2026)

Environmental Cost of AI's Energy Use:
Carbon, Water and Land Footprints
United Nations University Institute for Water,
Environment and Health
<https://digitallibrary.un.org/record/4116063?v=pdf>

Appendix

The following appendix table complements the main use-case overview in **Chapter 2** with additional examples that emerged from the workshop discussions and external expert presentations. The use cases are based in particular on contributions and presentations from PwC, Helaba, Novata, YUKKA and Syenah. They are included as supplementary illustrations of where and how

AI can support ESG-related banking processes, ranging from regulatory navigation and data enrichment to due diligence, monitoring and reporting. The table is not intended as a complete inventory, but as a practical catalogue of further options that banks may consider, adapt or use as inspiration when developing their own ESG-AI roadmap.

Table A1: Catalogue of additional AI Use Cases in ESG Management

ESG activity	Pain point	AI approach	Typical output	Value	Requirements
1. ESG regulatory chat / regulatory navigator (PwC, Novata)					
Interpretation of ESG regulation, supervisory expectations and internal policy requirements.	ESG regulation is complex, changes frequently and is difficult for business users to navigate efficiently.	Document retrieval, Q&A and summarisation on a curated regulatory knowledge base, ideally with links to original sources.	Plain-language answers, source references, regulatory summaries and targeted explanations for specific questions.	Faster access to regulatory knowledge, reduced dependency on manual searches and more consistent interpretation across teams.	Requires curated and version-controlled sources, expert validation, clear boundaries to avoid legal overreliance and documentation of answers used in processes.
2. CSRD / NFRD flagging of counterparties (PwC)					
Identification of counterparties subject to non-financial reporting obligations.	Banks often lack a complete and up-to-date overview of which counterparties publish relevant sustainability reports or are subject to reporting obligations.	AI-supported web search, report detection, template recognition and classification of reporting status.	Enriched counterparty list with reporting flags, identified report links or evidence, and confidence indicators.	Reduced manual research effort, faster portfolio enrichment and improved basis for reporting, risk analysis or client engagement.	Depends on reliable entity matching and transparent evidence. False positives or outdated web information require human review and audit trails.

ESG activity	Pain point	AI approach	Typical output	Value	Requirements
--------------	------------	-------------	----------------	-------	--------------

3. NACE code enrichment to level 4 (PwC)

Taxonomy reporting, credit portfolio analysis and regulatory disclosure.	Counterparty activity data is often only available at a coarse level, while regulatory analysis requires more granular classification.	Web search, activity extraction, mapping against NACE lists and plausibility checks against existing sector data.	Suggested four-digit NACE code, evidence from public sources and confidence score or exception flag.	Higher data granularity, improved data quality and more efficient preparation of taxonomy and portfolio analyses.	Requires clear classification logic, up-to-date NACE reference lists, review of ambiguous cases and safeguards against overconfident mapping.
--	--	---	--	---	---

4. KPI calculator and validator (PwC, Novata, Helaba)

Extraction and validation of ESG KPIs from sustainability, annual or taxonomy reports.	ESG KPIs are dispersed across reports, presented in different formats and difficult to validate consistently at scale.	Document retrieval, OCR, table extraction, KPI recognition and cross-checking against predefined definitions.	Structured KPI dataset with source locations, extracted values, validation flags and exceptions.	Faster data collection, improved consistency, reduced manual entry and better traceability for downstream reporting or portfolio analysis.	Requires robust table extraction, source references, data-quality controls and expert review where definitions or units differ.
--	--	---	--	--	---

5. Transition plan reader (PwC, Helaba)

Assessment of counterparties' transition plans and climate-related KPIs.	Transition plans are lengthy, heterogeneous and difficult to compare across clients or sectors.	AI-supported report search, KPI extraction, qualitative summarisation and comparison against predefined assessment criteria.	Summary of transition plan quality, extracted targets, progress indicators, sector alignment information and open gaps enriched with available internal information on the counterparty or contractual partner.	More focused credit and portfolio discussions, faster peer comparison and better preparation for transition-related client engagement.	Requires clear assessment criteria, sector expertise, careful handling of forward-looking statements and review of model interpretation.
--	---	--	---	--	--

ESG activity	Pain point	AI approach	Typical output	Value	Requirements
--------------	------------	-------------	----------------	-------	--------------

6. Qualitative taxonomy report drafting (PwC, Novata, Helaba)

Preparation of narrative parts of EU Taxonomy reporting.	Qualitative reporting involves repetitive drafting, consistency checks and integration of information from multiple sources.	Content generation based on structured inputs, prior reports, taxonomy templates and source-backed drafting workflows.	Draft narrative sections, suggested wording, consistency checks and lists of missing input.	Efficiency gains, more consistent language and reduced drafting burden for reporting teams.	Requires human approval, alignment with legal and disclosure standards, version control and safeguards against unsupported statements.
--	--	--	---	---	--

7. IRO generator and evaluator (PwC)

Support for double materiality assessment under CSRD.	Identifying and assessing impacts, risks and opportunities is resource-intensive and depends on multiple internal and external inputs.	Retrieval and synthesis of company information, sector data, regulatory expectations and easily adaptable assessment logic.	Suggested IRO longlist, preliminary scoring, rationale and evidence references.	More structured preparation of materiality workshops, improved completeness checks and comparability across business areas.	Requires methodological transparency, stakeholder validation and clear separation between AI-supported suggestions and final management judgement.
---	--	---	---	---	--

8. Automated taxonomy assessment / agentic workflow (PwC)

Taxonomy eligibility and alignment assessment in lending or periodic portfolio reviews.	Taxonomy assessment combines document collection, activity mapping, technical screening, DNSH, minimum safeguards and documentation, which makes the process complex and fragmented.	Agentic orchestration of specialised retrieval, extraction, mapping and checking steps, combined with human-in-the-loop approval.	Assessment package with extracted evidence, mapped activity, open questions, risk flags and documentation trail.	Potentially large efficiency gains, better process consistency and reusable components across taxonomy, reporting and credit workflows.	Requires strong governance, audit trails, escalation rules, access controls and careful consideration of EU AI Act implications where outputs influence decisions.
---	--	---	--	---	--

ESG activity	Pain point	AI approach	Typical output	Value	Requirements
--------------	------------	-------------	----------------	-------	--------------

9. Automated CSRD report update (PwC)

Sustainability reporting and annual report production.	Reporting teams must integrate changing regulation, internal data, prior-year content, management input and assurance requirements under time pressure.	Agentic or semi-agentic workflow combining source collection, regulatory scanning, drafting, XBRL preparation and quality checks.	Updated report sections, change logs, source lists, missing-data requests and consistency checks.	Reduced reporting workload, faster update cycles and improved quality control during reporting peaks.	Requires strong process ownership, source validation, approval workflow, version management and controls to avoid unsupported changes.
--	---	---	---	---	--

10. ESG and credit due diligence from adverse news (YUKKA, Syenah, Novata, Helaba)

Company-level due diligence for ESG, credit and investment risk.	Risk teams cannot manually screen millions of news items and may miss early signals of controversies or credit deterioration.	Entity recognition, metadata enrichment, Event Detection, Sentiment, Severity, query refinement and AI-generated due diligence reports with citations.	Severity-ranked due diligence report with risk events, supporting articles, citations and analytics.	Broader data coverage, faster investigation and earlier detection of relevant risk signals.	Requires transparent source citation, careful event clustering, entity disambiguation and review of false positives or noisy media signals.
--	---	--	--	---	---

11. Portfolio intelligence and incremental risk reporting (YUKKA, Syenah, Novata)

Ongoing portfolio monitoring for ESG, reputational and credit risks.	Portfolio managers need concise updates on what has changed since the last review rather than another full information dump.	Continuous news monitoring, severity ranking, memory function and customised report generation.	Daily or weekly incremental portfolio update, highlighting only material new developments and severe events.	Higher analyst productivity, sharper prioritisation and better early-warning capability across portfolios.	Requires clear definition of materiality, user-configurable relevance settings and governance over automated alert thresholds.
--	--	---	--	--	--

ESG activity	Pain point	AI approach	Typical output	Value	Requirements
--------------	------------	-------------	----------------	-------	--------------

12. Sector insights and market signals (YUKKA, Syenah)

Top-down sector monitoring for investment, credit and portfolio strategy.	Sector-level ESG and sentiment signals are difficult to derive manually from unstructured news at scale.	Aggregation of entity-level sentiment, ESG- and Risk- News signals into sector-level indicators and traffic-light views.	Sector dashboards, sentiment indicators, risk heatmaps and early-warning signals.	Improved market awareness, earlier identification of sector stress and additional context for portfolio steering.	Requires robust aggregation logic, transparent weighting, avoidance of overinterpretation and clear separation from investment recommendations.
---	--	--	---	---	---

13. Supply-chain and human-rights risk sentinel (Syenah, YUKKA, Novata)

Detection of hidden ESG risks in supply chains, including forced labour and sanctions-related exposure.	Tier-2 and Tier-3 suppliers are often invisible to financial institutions, while ESG signals may appear before they are captured in ratings.	Real-time monitoring of articles, NGO sources, regulatory lists, social media and relationship graphs across companies and supply chains.	Evidence-linked risk signal, supply-chain relationship view, severity assessment and audit trail.	Earlier detection of hidden exposure, better counterparty risk understanding and improved auditability of ESG risk monitoring.	Requires reliable relationship mapping, transparency of evidence chains, data-source quality controls and careful handling of indirect exposure.
---	--	---	---	--	--

14. Environmental incident early warning (YUKKA, Syenah, Novata, Helaba)

Monitoring of ecological incidents and operational ESG risks.	Environmental incidents can create reputational, operational and financial risks before traditional reporting or ratings reflect them.	Real-time event detection, severity scoring and evidence-based incident tracking across many source types and languages.	Incident alert, severity score, underlying evidence and follow-up indicators.	Faster reaction to ecological incidents, improved escalation and better integration of ESG events into risk surveillance.	Requires calibration of severity scores, filtering of irrelevant signals and clear escalation responsibilities.
---	--	--	---	---	---

15. Environmental and social due diligence cockpit for credit transactions (PwC)

E&S social due diligence in credit transactions before approval or escalation.	E&S checks are often manual, fragmented and difficult to document consistently across transaction workflows.	Multi-agent workflow with specialised agents, agent-to-agent revision loops and human-in-the-loop approval.	Due diligence assessment, policy-fit check, risk flags, escalation recommendation and audit trail for each transaction.	Faster decisions, reduced manual workload, more consistent checks and stronger traceability without replacing analyst judgement.	Requires clear decision ownership, role-based access, escalation logic, audit-ready documentation and controls for EU AI Act and model-risk considerations.
--	--	---	---	--	---

Approach of the Sustainable Finance Cluster Initiative

To assess the current state of AI in ESG Management in the banking sector, the Sustainable Finance Cluster launched an initiative in early 2026.

The following empirical methods were employed in the period of March - August to create insights and develop evidence-based recommendations. For further information about the research methods, feel free to contact the lead author.



IMAGE SOURCE: iStock.com/serayut, Bild-ID: 1181366400

Focus Group Workshops

Format: Pre-competitive exchange between institutions (presentations and discussions)

Focus: Key challenges, current practices and emerging solutions of AI Use Cases

Participants: Commerzbank, Deka, Deutsche Börse, DKB, DZ BANK, FIRM, Helaba, ING Germany, KfW, Novata

Survey

Focus: Implementation maturity, governance, value added, selection criteria and barriers & drivers.

Sample: the seven participating banks of the Cluster Initiative

Case Studies

Focus: Practical implementation insights from selected banks, covering the process, challenges and successes, outcomes and benefits, and next steps or remaining gaps.

Acknowledgements

This initiative has received funding from the Hessian Ministry of Economy, Energy, Transport, Housing and Rural Areas.

We thank the participants of the Cluster Initiative for their contributions:

Alexandra Musahl (Commerzbank)
Alicia Böning (DekaBank)
Christina Anselm (Novata)
Claudia Beck (ING)
Dr. Fabian Wendel (DZ BANK)
Katharina Pusch (Commerzbank)
Linus Vogel (Deka Investment)
Lukas Simon (Deutsche Börse)
Marina Preyssat (Helaba)
Marvin Hahn (Helaba)
Maryna Valiaieva (DekaBank)
Nam Nguyen (KfW)
Nico Schützhofer (KfW)
Oliver Hardel (DKB)
Sascha Seibert (DekaBank)
Dr. Sebastian Fritz-Morgenthal (FIRM)
Sebastian Moser (DZ BANK)
Sven Panz (ING)

We are also grateful for the contributions to the workshops from:

Andreas Pusch (YUKKA)
Kerim Bilican (PwC)
Mark Krause (Futury)
Philip Viertel (TechQuartier)
Selcuk Vural (PwC)
Steffen Müller (Salesforce)
Sven Gartung (Syenah)



IMAGE SOURCE: cl_stock_84081206 - stock.adobe.com

! Disclaimer

The views and opinions expressed in this position paper are those of the contributing authors and do not necessarily reflect the official policy or position of all stakeholders involved, including members of the focus groups.

This document solely serves as general background material in the field of AI in ESG Management. The authors and members of the Sustainable Finance Cluster have not specifically verified the information contained herein, nor can they be held responsible for any subsequent use which may be made of this information.

Claude (Anthropic) and Copilot (Microsoft) were used as assistants for text generation and text review. The ideas, analytical frameworks, structure and conclusions presented are entirely the authors' own. All AI-assisted work was critically reviewed and/or validated by the authors.

Imprint

Authors

Dr. Lukas Figge-Muschalik
lukas.figge@sustainablefinancecluster.com

Luca Steinhauer
Enzo Matos Balleux

Layout by YPSPUNKT Fotografie & Design

Sustainable Finance Cluster e.V.
www.sustainablefinancecluster.com

Managing Director

Matthias Hübner
matthias.huebner@sustainablefinancecluster.com