

Harnessing AI to Strengthen Supervisory Oversight of Nature-Related Disclosures

Chiara Colesanti Senni

University of Zurich

EBA-CETEx Workshop on Supervisory Approaches to Nature Related Risks
December 2025

The Challenge: Bridging Disclosure and Supervision

- Supervisors need to assess **public** nature disclosures at scale.
 - Beyond supervisory data, they are critical for market discipline, public accountability, and detecting greenwashing.
- Current disclosures frequently suffer from excessive length, complexity, and a lack of harmonisation (across regulation/guidance, jurisdictions, and formats).
- The problem is not lack of rules or guidance, but a gap between disclosure information and its efficient use for supervision.
 - Legal requirements (CSRD, CRR Pillar 3 disclosures) and soft law guidance (ECB Guide, EBA Guidelines) exist.
 - Supervisors need information that enable them to supervise EU banks and make sure they are financially sound (including nature-related risks).
 - E.g., CSRD focuses on disclosures of both impacts and risks for corporates, while Pillar 3 focuses on financial risk.
 - Supervisors need to cross-reference these “silos” to get a full picture of the risk, especially as Pillar 3 is not yet fully elaborated for nature.
- **What's needed:** Understand how to leverage available (but scattered) information for a science-informed and comprehensive supervision of nature-related risks.

- **Scale:** Analyze hundreds of reports systematically and consistently (same standards for each institution) across entire portfolios.
- **Detection:** Identify **gaps between talk and walk** and inconsistencies automatically across all filings (e.g., via sentiment analysis).
- **Depth:** Extract nuanced insights through sentiment analysis, geographic mapping, and cross-referencing (i.e., how meaningful disclosures are).
- **Efficiency:** Free up supervisory resources for high-risk cases and strategic oversight.
 - These tools can help supervisors more quickly identify which firms or areas require closer **(human) oversight**.

- A growing ecosystem of specialized AI tools is emerging to analyze disclosures. We can group them into:
 - **Climate-focused:** Tools for climate sentiment (ClimBERT) or transition plans.
 - **Integration-focused:** Tools for document processing and unstructured data.
 - **Nature-focused:** Tools like **AskNature**, which is our focus today.
- These nature tools are often aligned with the **TNFD framework**.
 - Leading **global benchmark** for nature-related risk management and disclosure.
 - Aligning the AI tool with TNFD provides a comprehensive, standardized, and widely-understood structure for the analysis.
- The key advantage of this AI approach is that the tools are **flexible and modular**.
 - This flexibility is crucial for a rapidly changing regulatory landscape.
 - E.g., New Pillar 3 requirements? You can update the tool.
 - Need EUDR or ESRS alignment? You can add that module.

Introducing AskNature: An AI-Powered Assessment Tool

- The first AI tool specifically built for nature disclosure assessment.
- **Core Architecture:** Retrieval-Augmented Generation (RAG).
 - It finds relevant passages first, then analyzes them.
- **Framework:** Aligned with the 14 TNFD requirements, covering all pillars (Governance, Strategy, Risk & Impact, Metrics & Targets).
- **Key Analyses:**
 - **Sentiment analysis (Two-way):**
 - Company-to-Nature: How do firms frame their impact on nature?
 - ← Nature-to-Company: How do they frame their dependence on nature?
 - **Geographic mapping:** Where do companies say they operate?
 - **Enforceability scoring:** Whether what they commit to is legally binding, i.e. whether a commitment is part of a formal regulatory filing or legal contract, vs. just a voluntary disclosure.
- **Application:** We applied this tool to the Nature Action 100 companies. The tool is publicly available to use and build on.

The process uses a Retrieval-Augmented Generation (RAG) architecture:

1. **Indexing:** All documents are first chunked and embedded in a vector database.
2. **Retrieval:** The tool uses semantic search to find the most relevant passages for each question.
 - E.g. A question about “biodiversity targets” will retrieve passages about “species conservation goals” even if the exact phrase differs.
3. **Tree of Questions:** The 14 TNFD questions are analyzed directly and broken down into 1,694 granular sub-questions for a deep, specific assessment.
 - E.g. “Do you have biodiversity targets?” becomes:
 - Are the targets quantitative?
 - Do they have specific timelines?
 - How is progress measured?
4. **Scoring:** The AI model runs 25 times per company to ensure reliability (as AI models are probabilistic).
 - The final output is a score plus detailed, sourced justifications for the assessment.

Key Finding: The “Talk vs. Walk” Problem (1/2)

- Positive Sentiment vs. Risk Reality:

- We find 72% of nature-related text has a positive sentiment, with companies framing themselves as “stewards” or “leaders.”
- This often reflects overcommunication on small, positive actions (like philanthropy) and not a realistic, risk-based analysis of their actual exposures.
- This suggests disclosures reflect aspirations, not reality, as global biodiversity indices continue to decline.

Key Finding: The “Talk vs. Walk” Problem (2/2)

- **The Four Pillars - An Imbalance:**

- Companies perform well on *narrative* pillars:
 - **Governance:** Describing board oversight, committees, policies.
 - **Strategy:** Building a narrative about the importance of nature.
 - **Risk & Impact:** Identifying high-level risks.
- They perform poorly on *actionable* pillars:
 - **Metrics & Targets:** Lacking quantitative targets, timelines, actual progress tracking, or risk mitigation measures.

- **Supervisory Implication:**

- One cannot rely on high scores in Governance and Strategy.
- Supervisors must examine whether quantitative metrics, progress tracking, and concrete mitigation measures exist.
- AI helps identify this pattern at scale.

Critical Gaps in Disclosures

- **Geographic and Supply Chain Gaps:**

- Some countries are underrepresented in disclosures.
- We find minimal overlap with known **biodiversity hotspots**.
- **Supply chains** remain largely invisible.
- Does this reflect a true lack of economic links, or is it a *failure of disclosure*? For many firms with global supply chains, the latter is likely.
- Under the upcoming EU Omnibus amendments (simplification agenda), banks cannot force smaller value-chain partners to report data for CSRD compliance.
- AI can infer supply chain locations and potential risks using proxy data, allowing supervisors to assess risk even without specific banks or non-financial corporates' disclosures.

- **The “Enforceability Gap”:**

- The scores show a clear pattern:
 - Avg. Commitment Score: 6.11 (Firms say they will do things)
 - Avg. Specificity Score: 5.29 (Firms set specific commitments)
 - Avg. Enforceability Score: 3.66 (Firms don't make it binding)
- If commitments are vague, non-specific, and not legally binding, supervisors **cannot assess the financial risk**.

- **Current State:**

- TNFD-aligned disclosures can *look* compliant on the surface.
- BUT they often lack the **decision-useful information** needed for risk assessment.
- While mandatory disclosure provisions are emerging (CSRD/ESRS, Pillar 3), the link to a risk-based assessment for nature is often weak or non-specific.

- **AI could enable:**

- Large-scale, automated screening of all institutions.
- Detection of “gaps between talk and walk” and boilerplate patterns.
- Systematically flagging where disclosures are too vague for a **risk-based assessment**.

- **Supervisory use cases:**

- Pre-assess disclosures against mandatory frameworks (CSRD/ESRS, Pillar 3).
- Conduct deeper due diligence on high-risk firms flagged by the AI.
- Benchmark performance across institutions and sectors.

Long Term:

- **The Framework:**

- Increase coherence between different parts of the disclosure framework (e.g. align the use of CSRD data points with Pillar 3 risk assessment needs).
- Supervisors can offer banks a range of options with varying intensity to gather essential risk data.

- **Obtain Data on Nature-related Risks:**

- Banks will not be able to require clients outside the CSRD scope to provide them CSRD-related information, but they can engage with clients to obtain the necessary information for risk management and prudential purposes.
- E.g. require geographic coordinates for key assets and operations.

- **Robust Verification:**

- Integrate remote sensing data (e.g., satellite imagery) to check claims.
- Track disclosed *activities* against real-world *outcomes* (greenwashing).

Short Term:

- Empower Supervision with AI Tools:

- Even without granular public disclosures, supervisors can already start using the tools to oversee nature-related risks
- **Deploy AI tools** to screen all disclosures for specificity, enforceability, and “gaps between talk and action.”
- Use AI-flagged reports to **prioritize resource-intensive human oversight** on the firms with the weakest, most evasive disclosures.

The Future: AI Agents & Adaptive Oversight

- **From Static Tools to Dynamic Agents:**
 - Next-generation AI agents can be designed to autonomously monitor, analyze, and flag issues in real-time as new disclosures are published.
- **Key Capabilities:**
 - Continuous monitoring of filings.
 - Cross-referencing disclosures with external data sources.
 - Adaptive learning from new regulations and supervisory feedback.
- **Integration with Digital Twins:**
 - AI agents can connect corporate disclosures with:
 - Satellite imagery and remote sensing data.
 - Biodiversity monitoring platforms.
 - Digital twin models of ecosystems.
 - This enables real-time verification (checking Disclosures vs. Reality).

Goal: A self-updating supervisory system that adapts to new frameworks

Key Takeaways

- The goal is to use available nature-related information to ensure financial safety.
- AI can help leverage available information for a science-informed supervision of nature-related risks
- The future of supervision is **adaptive**, combining AI agents with real-time verification (e.g., remote sensing) to check disclosures against reality.

Thank You

Chiara Colesanti Senni — University of Zurich

`chiara.colesantisenni@df.uzh.ch`
`surreal.ai`