doi: 10.15678/PG.2026.76.2.04

Journal of Public Governance
No. 2(76)/2026
ISSN 2956-6061 | e-ISSN 2956-6398

ACCEPTED MANUSCRIPT
Peer-reviewed and accepted for publication. Not yet copyedited, proofread or typeset.

Wioleta Gałat, Magdalena Zajączkowska

Mapping AI and ESG Best Practices: A Comparative Typology of European Project Cases

Abstract:

Objective: This article develops an original typology of artificial intelligence (AI) and environmental, social, and governance (ESG) best practices, examining how these two domains intersect in European organisational and educational contexts.

Research Design and Methods: Qualitative content analysis was applied to eleven best-practice cases collected across five countries (Estonia, France, Hungary, Slovenia, Poland) within the DEDYCATING Erasmus+ project, coded against seven dimensions.

Findings: AI is present in six of eleven cases but genuinely operational in only five: in one case AI is the subject of education rather than a tool in use, and the remaining five address substantive ESG objectives without current AI involvement. A 2×2 typology crossing AI implementation status with operational versus educational orientation yields four types: AI-Enabled ESG Operations, AI-Enabled ESG Pedagogy, ESG-Only Operations and Advocacy, and ESG-Only Education and Training.

Implications/Recommendations: "AI and ESG practice" functions as an umbrella label spanning fundamentally different configurations; future best-practice collections should require explicit AI-status disclosure and assess transferability against resource intensity and institutional dependency, not AI presence alone.

Contribution/Value Added: The study provides the first inductively grounded typology of AI-ESG best practices at the organisational and educational level in Europe, demonstrating that AI integration and ESG substance are largely independent design choices within this sample.

Article classification: research article

Keywords: artificial intelligence, ESG, best practices, typology, qualitative content analysis, higher education, sustainability

JEL classification: O33; Q56; I23; M14

About this version: This is the accepted manuscript of an article peer-reviewed and accepted for publication in the Journal of Public Governance. It may differ from the final version of record in language, formatting and pagination. The version of record will replace it on the journal website; the DOI will remain unchanged.

Wioleta Gałat *, PhD, Assistant Professor, Department of Public Management, Krakow University of Economics, Krakow, Poland, Rakowicka 27, Krakow, Poland; email: galatw@uek.krakow.pl; ORCID: 0000-0002-6807-8176
Magdalena Zajączkowska, PhD, Assistant Professor, Department of Public Management, Krakow University of Economics, Krakow, Poland, Rakowicka 27, Krakow, Poland; email: zajaczkm@uek.krakow.pl, ORCID: 0000-0001-7449-8698
* *corresponding author*

1. Introduction

Artificial intelligence (AI) and environmental, social, and governance (ESG) principles increasingly intersect in European business and higher education practice, driven by the near-simultaneous adoption of the EU Artificial Intelligence Act (European Parliament & Council of the European Union, 2024) and the Corporate Sustainability Reporting Directive (CSRD, 2022). Despite growing conceptual interest in the AI-ESG nexus, empirical evidence on how this intersection manifests in documented organisational and educational practice and how much of what is labelled "AI and ESG practice" genuinely involves AI remains scarce.

The practical case for integrating AI with ESG is well rehearsed. ESG data are dispersed and labour-intensive to collect, and the resulting disclosures have long been criticised as narrative and difficult to verify (Christensen et al., 2021); AI promises to automate data capture, standardise scoring, and enable verification at scale (De Villiers et al., 2024). The CSRD, the ISSB standards, and the EU AI Act make this promise more urgent, while also imposing new governance obligations on AI systems themselves. The relationship is therefore bidirectional: as AI becomes embedded in ESG processes, the risk that its capabilities are overstated, or "machinewashed", becomes itself an ESG concern (Seele & Schultz, 2022). Research on this nexus is growing rapidly but remains largely conceptual and pitched at the level of firms, sectors, and national systems (Vinuesa et al., 2020; Kar et al., 2022; Lee et al., 2025), while a largely separate strand treats generative-AI literacy in business and management education as a learning outcome in its own right. Four gaps cut across this literature.

First, grounded evidence of what AI-ESG integration actually looks like at the level of individual, documented practices is scarce, so integration is more often asserted than observed. Second, best-practice collections are widely used for benchmarking and lesson-drawing (Dolowitz & Marsh, 2000), yet the self-reporting bias they share with voluntary ESG disclosure (Christensen et al., 2021) is rarely quantified within the resulting analysis. Third, typology-building, an established form of theory construction well suited to small, heterogeneous samples (Doty & Glick, 1994; Fiss, 2011), has not been applied to ask whether the practices currently labelled "AI-ESG best practice" are better understood as qualitatively distinct configurations than as positions on a single maturity scale. Fourth, recent work on "machinewashing" warns that the AI label is frequently applied more widely than the underlying technology warrants (Seele & Schultz, 2022), yet this possibility has not been examined empirically within a collection explicitly framed around "AI and ESG."

This article addresses these gaps through a qualitative analysis of eleven best-practice cases collected across five countries (Estonia, France, Hungary, Slovenia, Poland) within the DEDYCATING Erasmus+ project.

The article's central contribution is an original, inductively derived typology of AI and ESG best practices. Rather than evaluating the eleven cases against a predefined maturity scale, this article asks a more basic classificatory question: what distinct types of AI-ESG practice exist within this collection, and along which dimensions do they differ? The analysis reveals a finding with consequences for how the field defines its own object of study: of the eleven cases submitted to a collection explicitly framed as documenting "AI and ESG best practices," fewer than half involve AI technology that is genuinely operational at the time of reporting. The remaining cases represent substantively legitimate, often highly transferable ESG practice in which AI plays no current role, AI-ethics education in which AI is the subject rather than the tool, or AI-literacy practice in which the ESG linkage is comparatively weak. This finding motivates the typology developed in Section 5, which differentiates AI-enabled operational practice, AI-enabled educational practice, and two categories of ESG practice in which AI is absent or merely planned.

The remainder of the article proceeds as follows. Section 2 reviews the literature on AI-ESG integration, best-practice research methodology, and typology-building in management and education research, situating the article's contribution relative to the broader literature reviewed there. Section 3 presents the methodology: the qualitative content analysis design, the coding scheme, and the procedure through which an additional "AI status" coding dimension emerged inductively from the data.

Four research questions guide the analysis:

RQ1: What is the sectoral distribution of the documented best practices, and which types of institutions are primarily responsible for their design and delivery?

RQ2: To what extent are the individual ESG pillars -Environmental (E), Social (S), and Governance (G) covered across the collected cases, and which dimensions dominate the current practice?

RQ3: Which specific classes of artificial intelligence technologies are deployed within the collection, and what is the actual operational implementation status of AI at the time of reporting?

RQ4: What qualitatively distinct configurations of AI-ESG practices can be inductively derived from the sample, and how does cluster membership relate to the self-reported transferability and resource intensity of the practices?

Section 4 reports the descriptive and cross-tabulated comparative findings (RQ1–RQ3). Section 5 develops the typology (RQ4). Section 6 discusses the findings in relation to the literature reviewed in Section 2, and Section 7 concludes with implications, limitations, and directions for future research.

2. Literature Review

2.1. AI-ESG Integration in Organisational and Educational Practice

The literature on AI-ESG integration remains nascent and dispersed across several disciplinary fields (finance, management, accounting, and information systems), with conceptual and normative contributions still outnumbering grounded empirical work (see also Lee et al., 2025; Sætra, 2021), particularly at the level of national policy and governance frameworks.

The broad premise animating this literature is that artificial intelligence can materially advance sustainability objectives: Vinuesa et al. (2020), in a widely cited expert-elicitation study, estimate that AI could act as an enabler for the majority of the targets underpinning the Sustainable Development Goals while simultaneously inhibiting a substantial minority, framing AI as a general-purpose capability whose sustainability value is contingent rather than automatic. Systematic reviews of the field (Kar et al., 2022) confirm both its rapid expansion and its persistent intellectual fragmentation, a structural feature that motivates the present article's attempt to observe AI-ESG integration at the level of concrete, documented practice rather than in the aggregate.

At the organisational level, the clearest locus of AI-ESG integration to date is ESG data management, disclosure, and assurance. Lee et al. (2025) propose a responsible-AI assessment framework explicitly linking AI investment decisions to ESG criteria, arguing that ESG-AI integration mitigates systemic risk and supports long-term value creation. This framework anticipates, at a conceptual level, the operational pattern observed empirically in Type I of the typology developed (AI deployed within ESG-relevant operational processes such as compliance reporting or environmental monitoring).

A parallel and more applied stream documents how natural-language processing, machine learning, robotic process automation, and digital-twin technologies are increasingly deployed to capture, score, and independently verify ESG information at scale, progressively reshaping ESG disclosure from a largely narrative, self-reported exercise into a more data-driven and externally verifiable one (De Villiers et al., 2024; Vilchez Olivares & Astorga De La Cruz,

2026). A growing body of firm-level empirical work reports, in turn, that AI adoption is positively associated with corporate ESG performance, typically through efficiency gains and improved resource allocation (e.g., Yu, Fan, & Yu, 2025).

Two features of this literature matter for the present article. First, Vilchez Olivares and Astorga De La Cruz (2026) classify studies by the depth of AI involvement, distinguishing cases in which AI is the methodological core from those in which it is merely an analytical support or a contextual reference, a graded logic that closely parallels the "AI status" distinction, at the level of individual practices. Second, this literature is framed predominantly at the level of firms, sectors, or national systems, and Lee et al.'s (2025) framework, although developed in consultation with industry, is normative rather than grounded in an empirical sample of documented practices. The present article complements that conceptual literature with grounded evidence of what such integration looks like, and how frequently it actually occurs, in a real collection of submitted best practices.

At the educational level, a distinct and more developed literature addresses the integration of generative AI into business and management education, largely separate from ESG content. Holmes et al. (2022) establish the ethical stakes of AI use in education generally, while a broad and fast-growing body of systematic-review evidence maps the competencies, benefits, and risks associated with generative-AI literacy across both educational and business settings (Reicho et al., 2026).

Two features of this literature are analytically important here. First, it consistently treats AI literacy as a multidimensional competency in its own right spanning technical, personal-interpersonal, and ethical-critical dimensions and therefore as a learning outcome distinct from any specific ESG content that AI might be used to teach or generate. Second, it confirms that the field remains predominantly conceptual rather than empirical, with only a minority of studies in either education or business grounded in empirical data (Reicho et al., 2026). This distinction between AI as subject matter and AI as a tool for delivering other subject matter becomes central in the analytical part, where several "AI and ESG" educational best practices are shown to be substantively AI-literacy practices with only a weak ESG-pillar linkage.

Cutting across both the organisational and educational strands is a growing concern that the "AI" label is applied more widely than the underlying technology warrants. Seele and Schultz (2022), reasoning by analogy from the established concept of greenwashing, define "machinewashing" as the practice of conveying misleading impressions of AI capability or of ethical AI use, and argue that the reputational incentives driving greenwashing operate with comparable force in the AI domain.

This concern is a direct counterpart, on the technology side, to the disclosure-quality problem long documented in voluntary and mandatory ESG reporting (Christensen et al., 2021): just as institutions have an incentive to present their sustainability credentials favourably, they have an incentive to foreground any association with AI, whether or not AI is operationally present. Notably, AI itself is increasingly used to detect the analogous phenomenon on the sustainability side, with natural-language-processing and AI-based tools deployed to assess the quality of firms' climate disclosures and their alignment with reporting frameworks such as the TCFD (Chabot, 2025), a symmetry that underscores how loosely the "AI" and "ESG" labels can attach to underlying substance.

In a best-practice collection assembled under the joint banner of "AI and ESG," these two labelling incentives compound. This literature therefore supplies the interpretive frame for one of the present article's principal empirical findings, that a substantial share of practices submitted as "AI and ESG best practice" involve no operational AI at all, and cautions against reading the mere presence of the "AI" label as evidence of genuine integration.

2.2. Best-Practice Collections as a Genre of Evidence

The empirical material analysed in this article belongs to a distinctive genre: the curated collection of self-reported "best practices". Such collections have a long history in management and public-policy research as vehicles for benchmarking, lesson-drawing, and the diffusion of successful practice across organisational and national boundaries (Dolowitz & Marsh, 2000).

Methodologically, they are typically approached through qualitative document analysis (Bowen, 2009) and content analysis (Krippendorff, 2018), which are well suited to the heterogeneous, free-text case descriptions examined here. Two features of the genre, however, condition what can be inferred from it.

First, best-practice collections are subject to a systematic selection effect: practices are nominated precisely because their authors regard them as successful, transferable, and impactful, so the resulting sample over-represents favourable cases relative to the underlying population of organisational practice. This self-reporting bias is well documented in the parallel literature on ESG disclosure, where the quality and neutrality of voluntarily disclosed information is a recurring concern (Christensen et al., 2021), and it applies with at least equal force to voluntarily submitted best-practice narratives.

Second, the literature on best-practice transfer cautions that the transferability a submitting institution claims for its practice is rarely straightforward to realise: internal "stickiness", arising from the recipient's absorptive capacity, causal ambiguity, and the resource intensity of the originating context, routinely impedes replication even of demonstrably effective practices (Szulanski, 1996). This body of work frames two of the present article's analytical choices: the decision to treat self-reported transferability as a claim to be examined rather than a measured property, and the finding, that resource intensity and institutional dependency, rather than the presence or absence of AI, are the primary constraints on transferability within the collection.

2.3. Typology-Building and Configurational Approaches

The article's central analytical contribution is a typology, and the choice of a typological rather than a scalar strategy is itself grounded in an established methodological literature. Typologies have been defended as a distinctive and rigorous form of theory-building rather than as mere classification.

Doty and Glick (1994) argue that a fully specified typology constitutes a set of ideal types whose relationships to outcomes can be stated and tested, and is therefore theoretically richer than a single ordinal ranking. This tradition has increasingly been articulated in configurational terms, in which cases are understood as combinations of co-occurring conditions rather than as points on a single continuum (Fiss, 2011; Ragin, 2014).

Configurational reasoning is particularly appropriate for the small-N comparative material analysed here. Ragin (2014) develops case-oriented comparative analysis for samples too small for conventional statistical treatment and too large for purely narrative comparison, and Fiss (2011) extends this configurational logic to typologies in organisation research; together, they underpin the constant-comparison procedure through which the present typology is derived.

The configurational perspective also carries a substantive implication that this article adopts directly: because the identified types are qualitatively distinct configurations rather than gradations of a single construct, they need not be rank-ordered, and a practice that lacks AI is not thereby a "less mature" version of one that deploys it. This position engages with methodological debates in the digital-government literature on how, and how far, development can be modelled as a sequence of stages (Solli-Sæther & Gottschalk, 2010); the present article carries that question from the level of national systems to the level of individual organisational and educational practices.

2.4. Positioning of the Present Article

The literature reviewed above identifies four gaps that this article addresses jointly. First, conceptual work on AI-ESG integration (Lee et al., 2025; Yu, Fan, & Yu, 2025) has not yet been tested against a grounded sample of documented organisational and educational practice. Second, the literature on voluntary disclosure (Christensen et al., 2021) flags a self-reporting bias that applies equally to best-practice collections but is rarely made explicit or quantified within the resulting analysis, rather than simply noted as a caveat. Third, the typology-building tradition (Doty & Glick, 1994; Fiss, 2011; Ragin, 2014) has not yet been applied at the level of individual AI-ESG practices to test whether a single "maturity" framing adequately captures the diversity of what is currently labelled AI-ESG best practice. Fourth, and cutting across the first three, the literature on machinewashing and disclosure quality (Christensen et al., 2021; Seele & Schultz, 2022) implies that a collection assembled under an "AI and ESG" banner may systematically overstate the presence of AI, yet this possibility is seldom examined empirically within a real collection.

This article addresses these gaps through three elements: a transparent coding methodology applied to a real best-practice collection, descriptive findings that explicitly quantify the self-reporting and AI-absence patterns in the data, and an original typology that treats AI-ESG practice as multi-categorical rather than unidimensional.

3. Methodology

This article adopts a qualitative content analysis design, consistent with the broader methodological tradition of comparative best-practice analysis. The analytical objects are not countries, institutions, or national education systems (a comparison already undertaken in prior research on national AI-ESG governance), but rather the individual best-practice cases themselves. The aim is to develop an original, inductively grounded typology of AI and ESG best practices, capturing recurring patterns in how artificial intelligence and ESG principles are combined in practice across business, higher education, and civil-society contexts, rather than to evaluate the comparative performance of the five participating countries.

This design choice follows directly from the nature of the data source. The Best Practice Collection comprises individually authored, free-text case descriptions submitted by partner institutions, each following a shared template but varying considerably in length, emphasis, and reporting style. Such material is well suited to qualitative coding aimed at typology-building, but poorly suited to the kind of standardised, point-scored maturity assessment applied to harmonised country-level reporting. Content analysis (Krippendorff, 2018) rather than indicator-based scoring is therefore the appropriate methodological response to this material.

As discussed in Section 2.2, best-practice collections of this kind are heterogeneous, self-submitted, and subject to a self-reporting bias analogous to the disclosure-quality problem in voluntary ESG reporting (Christensen et al., 2021). The present collection corroborates this skew empirically: eight of the eleven cases are rated as highly transferable, based substantially on each report's own self-authored assessment. This article treats that skew not as a flaw to be corrected but as a finding about best-practice collections as a genre of evidence, examined in the Discussion.

This article's 2×2 typology (AI status × operational/educational orientation) is derived through constant comparison rather than located on a single maturity scale. The consequence for this dataset is concrete: the four resulting types are qualitatively distinct rather than rank-ordered, so AI-absent cases such as the 25-year-old, institutionally embedded Forum Odpowiedzialnego Biznesu (BP9) or the empirically validated Sustainable HRM module (BP3), both reporting strong, durable outcomes, are not treated as merely "less mature" than AI-present ones.

The primary data source is the DEDYCATING Best Practice Collection DEDYCATING Project, 2025, comprising eleven best-practice reports submitted across the five participating countries (Estonia, France, Hungary, Slovenia, Poland) under Erasmus+ project 2024-2-HU01-KA220-HED-000278642. Ten reports were received, two per country; one of them (the Hungarian business report) presents two distinct cases - the Mercedes-Benz Full-Flex Plant in Kecskemét and the BMW iFACTORY plant in Debrecen - within a single report. Because these two cases differ in ownership structure (retrofit of an existing plant versus a greenfield project), in the specific AI technologies deployed (AI-assisted quality control and ESG data automation versus digital-twin simulation and predictive-maintenance AI), and in maturity stage (full production versus final construction phase), they were coded as two separate units of analysis (BP5a and BP5b) rather than as a single combined case. This yields a final analytic sample of N = 11 best-practice cases, summarised in Table 1.

Table 1. Overview of the best-practice source material

ID	Best practice	Country
BP1	University of Tartu – AI Ethics Ecosystem	Estonia
BP2	ESG Expert training programme (Swedbank / TalTech / KPMG Baltics)	Estonia
BP3	Sustainable Human Resource Management module	France
BP4	Building GenAI Literacy by Design	France
BP5a / BP5b	Mercedes-Benz Full-Flex Plant (Kecskemét) / BMW iFACTORY (Debrecen)	Hungary
BP6	Integrating AI tools into higher education: a framework	Hungary
BP7	SDG Toolbox and Methodology Handbook	Slovenia

BP8	Iskraemeco smart meter solutions (IE.7 RC)	Slovenia
BP9	Forum Odpowiedzialnego Biznesu (FOB)	Poland
BP10	CampusAI - Human+AI Collaboration platform	Poland

Source: own elaboration based on the DEDYCATING Best Practice Collection.

All 11 good practices were included in the coding process, in accordance with the procedure described below.

3.1. Coding Scheme

Each best-practice report was read in full and coded against six dimensions, derived from the categories established in the article's research design and following established procedures of qualitative document analysis (Bowen, 2009). The six dimensions - sector, ESG pillar, AI technology, main objective, expected benefits, and transferability - are defined operationally in Table 2.

Table 2. Coding scheme: dimensions and operational definitions

Dimension	Categories	Operational definition
Sector	Public administration / Higher education / Business / NGO / Cross-sector	The institutional type primarily responsible for designing and delivering the practice, as self-identified in the country / institution field of each best-practice report.
ESG pillar	Environmental (E) / Social (S) / Governance (G) / combinations	The substantive ESG dimension(s) addressed by the practice's content or impact, coded as single or multiple pillars where the practice is explicitly cross-cutting; each pillar code was further classified as primary or secondary.
AI technology	Generative AI / NLP / Machine Learning / Predictive Analytics / Automation–RPA / Recommendation systems / Digital twin / Embedded computing–IoT / Other / None	The specific class of AI (or AI-adjacent digital) technology referenced as actually used in the practice, distinct from technologies merely mentioned as aspirational or planned.
Main objective	ESG reporting / Education & training / Decision support / Environmental monitoring / Citizen engagement / Compliance & governance / Operational efficiency	The primary functional purpose the practice is designed to serve, coded as the single best-fitting category; a secondary objective was recorded where explicitly dual-purpose.
Expected benefits	Transparency / Efficiency / Cost reduction / Inclusion / Sustainability / Better decision-making / Knowledge development	The outcome(s) the practice report claims or evidences, coded from the "Success factors," "Outcomes and impact," and "Justification" sections of each report.
Transferability	Low / Medium / High	The coder's assessment of how readily the practice could be replicated in a different national or institutional context, based on the explicit "Transferability / Adaptability" section of each report (necessary conditions, resource intensity, dependence on context-specific actors).
AI status	Implemented / Subject matter only / Absent (incl. planned only)	An additional, inductively derived dimension distinguishing practices in which AI is genuinely operational at the time of reporting from those in which AI is absent from current practice or referenced only as a future development plan. "Subject matter only" denotes practices in which AI is the topic (e.g., AI-ethics education) but no AI technology is used; for the typology, such cases are grouped with AI-absent cases, because AI is not operational.

Source: own elaboration.

Three of the six dimensions (sector, ESG pillar, main objective) permitted multiple codes per case where a report was explicitly cross-cutting; for example, the Mercedes-Benz case (BP5a) was coded as addressing all three ESG pillars, with Environmental as the primary

pillar, reflecting its combination of decarbonisation (E), workforce retraining (S), and ESG-data governance (G). For the ESG pillar, each code was classified as primary (a pillar the practice mainly addresses) or secondary (explicitly addressed but subordinate); cases with a broad, balanced ESG scope were coded as primary on all three pillars. The remaining three dimensions (AI technology, expected benefits, transferability) were coded as the single best-fitting category, or as an explicit combination where the report stated more than one benefit or technology with comparable weight.

3.2. Emergence of the "AI Status" Dimension

During the initial coding pass, a pattern emerged that the original six-dimension scheme did not anticipate: several reports submitted to a collection explicitly framed as "best practices in AI and ESG" described practices in which artificial intelligence was not, in fact, part of current operations. In some cases AI was absent from the practice altogether (e.g., the Estonian ESG Expert training programme, BP2, and the Polish Forum Odpowiedzialnego Biznesu, BP9, both of which are delivered entirely through human expertise, publications, and stakeholder dialogue); in one case, AI was the subject of the practice rather than a technology used within it (the University of Tartu AI Ethics Ecosystem, BP1, which consists of seminars, workshops, and a national collaboration network on AI ethics); in others, AI was explicitly named only as a future development plan rather than a current feature (e.g., the IEDC–Bled SDG Toolbox, BP7, described as "currently rooted in document-based dissemination," with AI-supported content curation listed under "future development plans"; and the Iskraemeco smart meter case, BP8, in which the authors state directly that "the technology is mainly hardware and embedded computing," with AI appearing only via the company's participation in an unrelated external research workshop).

This finding could not be accommodated within the existing "AI technology" category alone, because that category was designed to classify which type of AI is used, not whether AI is used at all. A seventh coding dimension - AI status, coded as Implemented, Subject matter only, or Absent (including planned-only) - was therefore added inductively after the first coding pass and applied retrospectively to all eleven cases. This departure from a strictly deductive, pre-registered coding scheme is consistent with standard practice in qualitative content analysis, where categories may be "data-driven" as well as "concept-driven" (Krippendorff, 2018), and is reported transparently. The decision to retain and report this dimension, rather than treat it as a coding inconvenience, follows directly from the article's research questions: RQ3 ("which AI technologies dominate?") cannot be meaningfully answered without first establishing for how many cases AI is actually present to be characterised.

4. Comparative Analysis of Best Practices

This section reports the results of the coding exercise, addressing RQ1–RQ3.

4.1. Sectoral Distribution and ESG Pillar Coverage (RQ1–RQ2)

Higher Education accounts for the largest share of cases (n=5, five of eleven), followed by Business (n=4), with one cross-sector business–HEI case (BP2) and one NGO-led case (BP9) (Table 3). Notably, Public administration, one of the five categories in the original coding scheme, is entirely absent from the collection, which limits the generalisability of any conclusions about government-sector AI-ESG practice.

Table 3. Distribution of best practices by sector

Sector	n	Cases (ID)
Business	4	BP5a, BP5b, BP8, BP10
Higher Education	5	BP1, BP3, BP4, BP6, BP7
Cross-sector (Business + HEI)	1	BP2
NGO / Cross-sector	1	BP9

Source: own elaboration based on coding sheet.

Table 4 reports ESG pillar coverage, distinguishing the pillars a practice primarily addresses from those it addresses in a subordinate way. Coverage is broad: each pillar is addressed by at least nine of the eleven cases, and four cases (BP2, BP6, BP7, BP9) address all three pillars in a balanced way. Differences emerge in primary focus. Environmental and Governance are each a primary pillar in seven cases, whereas Social is a primary pillar in five, and only one case (BP3) is primarily a Social practice. A single-pillar Environmental focus is confined to the industrial and energy-technology cases (BP5a, BP5b, BP8), while a single-pillar Governance focus characterises the AI-literacy and AI-ethics cases (BP1, BP4, BP10).

Table 4. ESG pillar coverage across best practices

ESG pillar	Primary / secondary (n)	Cases (ID)
Environmental (E)	7 / 2	Primary: BP2, BP5a, BP5b, BP6, BP7, BP8, BP9; secondary: BP3, BP4
Governance (G)	7 / 4	Primary: BP1, BP2, BP4, BP6, BP7, BP9, BP10; secondary: BP3, BP5a, BP5b, BP8
Social (S)	5 / 4	Primary: BP2, BP3, BP6, BP7, BP9; secondary: BP1, BP5a, BP5b, BP8

Source: own elaboration.

Primary = a pillar the practice mainly addresses; secondary = explicitly addressed but subordinate. Cases with a broad, balanced ESG scope (BP2, BP6, BP7, BP9) are coded as primary on all three pillars; a single case may address multiple pillars.

These pillar distributions reflect the composition of the sample rather than a claim about ESG priorities more broadly; as shown in Section 5, the presence of specific pillars is closely associated with sector type and AI implementation status.

4.2. AI Technologies, Implementation Status and Transferability (RQ3)

Of the eleven cases, only five involve AI technology that is genuinely operational at the time of reporting; the remaining six reference AI as absent or as a stated future plan. One of these six cases (BP1) nonetheless engages with AI as subject matter, through education on AI ethics, without using AI technology; AI is thus present in six of the eleven cases, but operational in only five. Table 5 cross-tabulates AI status against sector. Business cases show the highest implementation rate (three of four), Higher Education follows (two of five), and the NGO and cross-sector cases show none. This sectoral gradient suggests a structural rather than incidental pattern: AI-ESG integration is currently most mature in resource-intensive industrial and platform-business contexts.

Table 5. AI implementation status by sector

Sector	AI implemented	AI absent / planned	Total
Business	3	1	4
Higher Education	2	3	5
Cross-sector (Business+HEI)	0	1	1
NGO / Cross-sector	0	1	1
Total	5	6	11

Source: own elaboration based on coding sheet.

Among the five AI-implemented cases (Table 6), Generative AI is the most common technology (three of five implemented cases), spanning content creation (BP6), pedagogically embedded GenAI literacy (BP4), and commercial multi-model platforms (BP10). The remaining two cases each deploy a distinct technology: predictive analytics / AI-based quality control (BP5a) and digital twin combined with predictive-maintenance AI (BP5b).

Table 6. AI technology distribution among AI-implemented cases (n=5)

AI technology	n	Cases (ID)
Generative AI (incl. multi-model platforms)	3	BP4, BP6, BP10
Predictive analytics / AI-based quality control	1	BP5a
Digital twin	1	BP5b

Source: own elaboration. Counts refer to the AI-implemented subset (n=5).

Eight of the eleven cases are coded as High transferability and three as Medium (Table 7). This skew is treated with caution: each report's transferability assessment is self-authored by the submitting institution, introducing a well-documented disclosure-quality bias (Christensen

et al., 2021). The only Medium-transferability case within the AI-implemented subset (BP5b) involves capital-intensive infrastructure, while among the AI-absent cases BP2 depends on a context-specific tripartite institutional arrangement, suggesting that transferability is constrained by resource intensity and institutional dependency rather than by AI status itself.

Table 7. Transferability by AI implementation status

AI status	High transferability	Medium transferability	Total
AI implemented	4	1	5
AI absent / planned	4	2	6
Total	8	3	11

Source: own elaboration based on coding sheet.

Together, Tables 5 through 7 establish that AI implementation status and transferability are largely independent of one another within this sample, a finding whose implications are developed in the typology.

5. Typology of AI and ESG Best Practices

Building on the cross-tabulation of AI status (Table 5) and main objective, this section develops a 2×2 typology addressing RQ4. Two dimensions emerged as jointly sufficient to differentiate the eleven cases: (i) AI implementation status (Implemented vs. Absent/Planned) and (ii) orientation (operational/compliance vs. educational/capacity-building). Table 8 presents the typology.

Table 8. A 2×2 typology of AI and ESG best practices

	Operational / Compliance Orientation	**Educational / Capacity-Building Orientation**
AI Implemented	Type I AI-Enabled ESG Operations (BP5a, BP5b)	Type II AI-Enabled ESG Pedagogy (BP4, BP6, BP10)
AI Absent / Planned	Type III ESG-Only Operations and Advocacy (BP8, BP9)	Type IV ESG-Only Education and Training (BP1, BP2, BP3, BP7)

Source: own elaboration based on cross-tabulation of AI status (Table 5) and main objective.

Table 9 summarises the principal characteristics of each type before they are characterised in the paragraphs below.

Table 9. Summary characteristics of the four typology types

Characteristic	Type I	Type II	Type III	Type IV
n (cases)	2	3	2	4
Dominant sector	Business (100%)	HEI / EdTech (100%)	Business / NGO	HEI / Cross-sector (100%)
Dominant ESG pillar	E (S, G secondary)	G (BP4, BP10); broad (BP6)	E, S, G (broad)	Broad (BP2, BP7); S (BP3); G (BP1)
Primary objective	Compliance; efficiency	Education & training	Monitoring; advocacy	Education & training

Typical transfera-bility	High–Medium	High	High	High–Medium
Resource intensity	High (capital)	Low–Medium	Low–Medium	Low

Source: own elaboration.

Type I practices embed AI into core ESG-relevant business processes: AI-assisted quality control and SAP-based ESG reporting (BP5a, Mercedes-Benz) and digital-twin-enabled plant design with AI-supported predictive maintenance (BP5b, BMW iFACTORY). Both are primarily Environmental practices, with Social and Governance as secondary pillars, and report Compliance & governance or Operational efficiency as the primary objective. Transferability varies within the type: the retrofit model of BP5a is described as highly transferable, whereas the greenfield model of BP5b requires capital investment and infrastructure that constrain replication outside a comparable manufacturing context. The key finding is that AI-enabled and highly transferable are not synonymous: transferability within Type I depends on the scale of the surrounding operational context, not the AI component itself.

Type II practices use generative, frequently multi-model AI as a teaching tool for ESG- or sustainability-adjacent competencies. All three cases (ICN's GenAI Literacy by Design, BUEB's AI-in-HE framework, and CampusAI) report High transferability and Education & training as the primary objective. Critically, Type II is the type with the weakest direct ESG-pillar linkage: BP4 frames AI literacy as a Governance-adjacent competency rather than ESG content; BP6 is a partial exception, as its sustainability courses cover all three pillars, but AI serves only to generate course materials and is not itself linked to ESG content (the practice was also at the pilot stage at the time of reporting); and BP10's own authors describe it as "not strictly an ESG- or ethics-related project." With the partial exception of BP6, Type II practices are better understood as AI-literacy practices situated adjacent to ESG education than as ESG practices that happen to use AI.

Type III practices pursue substantive Environmental, Social, and Governance objectives without current AI involvement. Iskraemeco's smart meter technology (BP8) operates on hardware and embedded computing, its own authors state this explicitly, and FOB (BP9) relies on reports, surveys, and multi-stakeholder dialogue developed over 25 years. Both are coded as High transferability, illustrating that mature, impactful ESG practice does not require AI as a precondition. Both also flag AI as a future direction, positioning Type III as transitional rather than terminal.

Type IV comprises educational practices without operational AI. The Estonian ESG Expert training programme (BP2), ICN's Sustainable HRM module (BP3, which explicitly dis-

courages AI reliance in assessment), and IEDC-Bled's SDG Toolbox (BP7) represent the collection's most ESG-substantive educational cases: BP2 and BP7 address all three pillars, and BP3 is the only case in the collection whose primary focus is the Social pillar. The fourth case, the University of Tartu AI Ethics Ecosystem (BP1), is a distinctive variant: AI is the subject of the practice rather than a technology used within it, and its ESG linkage runs mainly through the Governance pillar (ethical and regulatory oversight of AI). Type IV stands in instructive contrast to Type II: where Type II is AI-rich but ESG-light, the core of Type IV is ESG-substantive but AI-absent, confirming that AI integration and ESG substance are largely independent design choices within this collection.

Within this sample, the typology's two axes appear independent: all four quadrants are populated, and sector membership closely tracks type membership (Business - Type I or III; HEI/EdTech - Type II or IV), although the small number of cases per type (two to four) precludes any stronger claim. The central implication is that "AI and ESG best practice" functions as an umbrella label spanning practices with fundamentally different relationships to AI technology, which any future collection instrument or evaluation rubric should make explicit.

6. Discussion

The finding that six of eleven cases involve no operational AI, combined with the ESG-light profile of Type II, indicates that "AI and ESG" is not a single coherent category. Two interpretations are possible. The first is collection-specific: partner institutions defaulted to whatever prominent initiative was available, because genuinely AI-integrated ESG practice is still rare. The second is structural: AI-ESG integration and substantive ESG practice are currently separate communities of practice within most organisations, pursued by different teams with limited overlap. The sectoral gradient in Table 5 (three of four in Business versus none in NGO/cross-sector cases) is consistent with the structural interpretation. In either case, the implication for future collections is clear: a collection instrument that implicitly rewards AI presence as a marker of sophistication risks systematically under-representing the substantively strong, non-AI-dependent ESG practice exemplified by Types III and IV. The pattern also gives empirical substance to the concern that the AI label attaches more widely than the underlying technology warrants (Seele & Schultz, 2022), although here the overstatement arises at the level of the collection's framing rather than in the individual reports, most of which state openly that AI is absent or only planned. The AI-ethics case (BP1) adds a further nuance: AI can enter a practice as subject matter rather than as a tool, echoing the distinction

drawn in the literature on AI literacy (Reicho et al., 2026). Counting such cases as AI practice would raise the share to six of eleven, but would conflate teaching about AI with using it.

Transferability is driven primarily by resource intensity and institutional dependency, not by whether a practice involves AI. This is consistent with the literature on best-practice transfer, which locates the main barriers to replication in the resources and absorptive capacity that a practice requires (Szulanski, 1996). AI-enabled practices are not systematically less transferable than non-AI ones; reduced transferability is explained by capital intensity (BP5b, an AI-enabled case) or by reliance on a non-substitutable tripartite institutional configuration (BP2, an AI-absent case), not by the AI component itself. At the same time, eight of the eleven cases rate their own transferability as High, a skew consistent with the selection effect typical of self-submitted best-practice collections and with the disclosure-quality problem documented for voluntary ESG reporting (Christensen et al., 2021); within this genre, self-reported transferability is better read as a claim than as a measured property. Two practical implications follow for future collection design: first, submission templates should require a binary AI-status disclosure rather than relying on post hoc inference; second, self-reported transferability assessments should be complemented by an independent checklist addressing resource intensity and institutional dependency.

The typology can also serve as a practical diagnostic tool for institutions. A university or company can map its current initiatives onto the four quadrants by asking two simple questions: is AI operational in the practice today, and is the practice mainly operational or educational? Such a mapping shows where the institution's AI-ESG portfolio is concentrated and where it is thin. Each quadrant then suggests a different development path rather than a single upgrade route. For Type II practices, the priority is to strengthen explicit ESG content, for example by linking GenAI-literacy courses to concrete E, S, or G learning outcomes, as the Sustainable HRM module (BP3) does. For Type III and IV practices, AI should be introduced only where it solves a specific problem, such as automating data collection for sustainability reporting, and not as a signal of modernity. For Type I practices, the main task is governance: documenting AI systems in line with the AI Act and disclosing their role transparently in CSRD reporting, which reduces the risk of machinewashing. Finally, funders and project consortia, including Erasmus+ partnerships, could use the typology when designing calls and evaluation criteria, asking applicants to state explicitly which configuration their practice represents.

7. Conclusions

This article has developed an original, inductively grounded typology of AI and ESG best practices, based on qualitative content analysis of eleven cases submitted across five European countries within the DEDYCATING project. The typology crosses two dimensions that appear independent within this sample - AI implementation status and operational/educational orientation yielding four types: AI-Enabled ESG Operations (Type I), AI-Enabled ESG Pedagogy (Type II), ESG-Only Operations and Advocacy (Type III), and ESG-Only Education and Training (Type IV).

The article's central empirical finding is that AI technology is genuinely operational in only five of the eleven cases submitted to a collection explicitly framed as documenting AI and ESG best practice; the remaining six cases represent substantively legitimate ESG or AI-ethics practice in which AI technology is either entirely absent or referenced only as a future development plan. This finding, which the original six-dimension coding scheme did not anticipate and which emerged inductively during coding, carries a normative implication for the field: "AI and ESG best practice" functions, at present, as an umbrella label spanning at least two qualitatively distinct phenomena - organisations and institutions that have genuinely integrated AI into ESG-relevant processes, and organisations and institutions pursuing substantive, often highly transferable ESG work that has not yet, or may never need to, incorporate AI technology as a precondition for impact.

From a practical standpoint, the findings carry implications for three audiences. For institutions designing future best-practice collections, the recommendation developed is to require explicit, binary AI-status disclosure at the point of submission rather than relying on post hoc inference from free-text description, and to complement self-reported transferability claims with an independent checklist addressing resource intensity and institutional dependency, the two characteristics this article identifies as the strongest correlates of reduced transferability, in contrast to AI status itself, which shows little independent association with transferability.

For institutions seeking to develop their own AI-ESG practice, the typology offers four distinct, equally legitimate models to draw upon, rather than a single "maturity" target to climb toward: an organisation strong in substantive ESG practice but without AI capacity (Type III/IV) is not thereby a less advanced version of an AI-enabled organisation (Type I/II), but is pursuing a different, viable configuration of ESG work.

For researchers studying AI-ESG integration, the article's central finding that AI presence and ESG substance are largely independent design choices within this sample, co-occurring in only a minority of cases suggests that future studies should resist treating "AI and ESG" as

a single coherent construct and should instead specify which of the typology's four configurations, or their analogues in other samples, a given study or policy instrument is actually addressing.

This study is subject to several limitations that should inform the interpretation of its findings. The eleven cases were not randomly sampled but were submitted according to an internal project quota of two best practices per country, and the sample's composition reflects partner institutions' own judgement of what constitutes a notable practice; the typology developed here should accordingly be understood as a grounded description of patterns within this specific collection, not as a claim to representativeness of AI-ESG practice across Europe more broadly. The apparent skew toward High transferability (eight of the eleven cases) is most plausibly an artefact of self-reported, favourably framed source material rather than an independently verified property of the practices themselves, and the article has treated this skew as a substantive finding about best-practice collections as a genre of evidence rather than attempting to correct for it analytically. The coding was carried out by a single coder without independent double-coding, so inter-coder reliability could not be assessed. Finally, the complete absence of a Public administration case in this collection means the typology's applicability to public-sector AI-ESG practice specifically remains untested.

Future research should pursue at least three directions building directly on these limitations. First, the typology developed here should be tested against a larger, ideally more representative or randomly sampled collection of AI-ESG practices, including public-sector cases, to assess whether the four types identified here and particularly the empirical independence of AI status from operational/educational orientation replicate outside this specific project's collection. Second, future coding of best-practice collections of this kind should incorporate independent double-coding with reported inter-coder reliability statistics, addressing the single-coder limitation acknowledged here. Third, the provisional finding that resource intensity and institutional dependency, rather than AI status, are the strongest correlates of transferability merits direct testing through a dedicated transferability instrument applied prospectively to new best-practice submissions, rather than inferred retrospectively from self-authored transferability narratives as in the present study.

References

Bowen, G. A. (2009). Document analysis as a qualitative research method. Qualitative Research Journal, 9(2), 27–40. https://doi.org/10.3316/qrj0902027

Chabot, M. (2025). Textual and AI-based analysis of climate disclosures: Evidence from the European energy sector. The British Accounting Review, Article 101736. https://doi.org/10.1016/j.bar.2025.101736

Christensen, H. B., Hail, L., & Leuz, C. (2021). Mandatory CSR and sustainability reporting: Economic analysis and literature review. Review of Accounting Studies, 26(3), 1176–1248. https://doi.org/10.1007/s11142-021-09609-5

Corporate Sustainability Reporting Directive (CSRD). (2022). Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 as regards corporate sustainability reporting. Official Journal of the European Union, L 322, 15–80.

DEDYCATING Project. (2025). Best practice collection in AI and ESG (Erasmus+ project 2024-2-HU01-KA220-HED-000278642). https://dedycating-project.hu/best-practice-collection/

De Villiers, C., Dimes, R., & Molinari, M. (2024). How will AI text generation and processing impact sustainability reporting? Critical analysis, a conceptual framework and avenues for future research. Sustainability Accounting, Management and Policy Journal, 15(1), 96–118. https://doi.org/10.1108/SAMPJ-02-2023-0097

Dolowitz, D. P., & Marsh, D. (2000). Learning from abroad: The role of policy transfer in contemporary policy-making. Governance, 13(1), 5–23. https://doi.org/10.1111/0952-1895.00121

Doty, D. H., & Glick, W. H. (1994). Typologies as a unique form of theory building: Toward improved understanding and modeling. Academy of Management Review, 19(2), 230–251. https://doi.org/10.5465/amr.1994.9410210748

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L, 2024/1689.

Fiss, P. C. (2011). Building better causal theories: A fuzzy set approach to typologies in organization research. Academy of Management Journal, 54(2), 393–420. https://doi.org/10.5465/amj.2011.60263120

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Buckingham Shum, S., Santos, O. C., Rodrigo, M. T., Cukurova, M., Bittencourt, I. I., & Koedinger, K. R. (2022). Ethics of AI in education: Towards a community-wide agenda. International Journal of Artificial Intelligence in Education, 32(3), 504–526. https://doi.org/10.1007/s40593-021-00239-1

Kar, A. K., Choudhary, S. K., & Singh, V. K. (2022). How can artificial intelligence impact sustainability: A systematic literature review. Journal of Cleaner Production, 376, 134120. https://doi.org/10.1016/j.jclepro.2022.134120

Krippendorff, K. (2018). Content analysis: An introduction to its methodology (4th ed.). SAGE.

Lee, S. U., Perera, H., Liu, Y., Xia, B., Lu, Q., Zhu, L., Cairns, J., & Nottage, M. (2025). Integrating ESG and AI: A comprehensive responsible AI assessment framework. AI and Ethics, 5(5), 5121–5148. https://doi.org/10.1007/s43681-025-00741-5

Ragin, C. C. (2014). The comparative method: Moving beyond qualitative and quantitative strategies. University of California Press.

Reicho, M., Otrel-Cass, K., Ebner, M., Brünner, B., Freinhofer, D., & Irfan, J. (2026). Generative AI literacy across education and business: Competencies, obstacles, and benefits - a systematic literature review. International Journal of Educational Technology in Higher Education, 23(1), Article 23. https://doi.org/10.1186/s41239-026-00596-8

Sætra, H. S. (2021). A framework for evaluating and disclosing the ESG-related impacts of AI with the SDGs. Sustainability, 13(15), Article 8503. https://doi.org/10.3390/su13158503

Seele, P., & Schultz, M. D. (2022). From greenwashing to machinewashing: A model and future directions derived from reasoning by analogy. Journal of Business Ethics, 178(4), 1063–1089. https://doi.org/10.1007/s10551-022-05054-9

Solli-Sæther, H., & Gottschalk, P. (2010). The modeling process for stage models. Journal of Organizational Computing and Electronic Commerce, 20(3), 279–293. https://doi.org/10.1080/10919392.2010.494535

Szulanski, G. (1996). Exploring internal stickiness: Impediments to the transfer of best practice within the firm. Strategic Management Journal, 17(S2), 27–43. https://doi.org/10.1002/smj.4250171105

Vilchez Olivares, P. A., & Astorga De La Cruz, B. J. (2026). AI and data analytics in sustainable financial reporting and ESG disclosure: A systematic literature review. Sustainability, 18(11), Article 5393. https://doi.org/10.3390/su18115393

Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., & Fuso Nerini, F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. Nature Communications, 11, Article 233. https://doi.org/10.1038/s41467-019-14108-y

Yu, X., Fan, L., & Yu, Y. (2025). Artificial intelligence and corporate ESG performance: A mechanism analysis based on corporate efficiency and external environment. Sustainability, 17(9), 3819. https://doi.org/10.3390/su17093819

Funding
This work has been supported by:
ERASMUS+ KA220-HED - Cooperation Partnerships
2024-2-HU01-KA220-HED-000278642
DEDYCATING - Development of dynamically changing curricula in ESG education supported by generative Artificial Intelligence
and
DEDYCATING – Dinamikusan változó tantervek fejlesztése az ESG-oktatásban, generatív Mesterséges Intelligencia támogatásával

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.



Co-funded by
the European Union



Research Ethics Committee
Not applicable.
Conflicts of Interest
The authors declare no conflicts of interest.
Copyright and License
This article is published under the terms of the Creative Commons Attribution 4.0 License. Published by Małopolska School of Public Administration – Krakow University of Economics, Krakow, Poland.
Data Availability Statement
The data supporting the findings of this study (five harmonised country reports: Poland, Hungary, Estonia, Slovenia, France) are available from the corresponding author upon reasonable request.
Author Contributions
W.G.: conceptualisation, methodology, writing – original draft, writing – review and editing; M.Z.: literature review, writing – original draft, review and editing. All authors have read and agreed to the published version of the manuscript.
AI Declaration Statement
The authors used generative AI tools for language-related assistance, including translation and minor editorial improvements during the preparation of the manuscript. All AI-assisted outputs were reviewed and verified by the authors, who retain full responsibility for the accuracy, interpretation, and final content of the article.