



Mapping sustainability in higher education: a natural language processing analysis of student theses

Downloaded from: <https://research.chalmers.se>, 2026-08-31 15:56 UTC

Citation for the original published paper (version of record):

Sundemo, M., Löfgren, Å., Yu, Y. et al (2026). Mapping sustainability in higher education: a natural language processing analysis of student theses. Higher Education, 92(1): 121-143. <http://dx.doi.org/10.1007/s10734-025-01526-9>

N.B. When citing this work, cite the original published paper.



Mapping sustainability in higher education: a natural language processing analysis of student theses

Mattias Sundemo¹ · Åsa Löfgren¹ · Yinan Yu² · Samuel Scheidegger³

/ Accepted: 10 August 2025 / Published online: 30 August 2025
© The Author(s) 2025

Abstract

This paper examines the integration of sustainability in higher education, focusing on how business, economics, and finance students relate their work to the Sustainable Development Goals (SDGs). We develop SustainSight, an inference engine using natural language processing to assess sustainability content in written text. Using SustainSight, we analyze over 23,000 bachelor's and master's theses from the five largest Swedish universities offering programs in business, economics, and finance spanning a 15-year period. This unique dataset allows us to provide insights into how sustainability is reflected in students' theses, with a focus on variation by gender, education level (bachelor/master), discipline, and university. For the subset of theses, for which we have the authors educational program and data on curriculum reforms, a staggered difference-in-differences model is estimated. We find that integration of sustainability into degree requirement substantially increases the likelihood that students include sustainability perspectives into their theses. Finally, we conduct a topic analysis to explore in greater depth what students write about when addressing sustainability-related issues.

Keywords NLP · Sustainability assessment · Higher education · Business and economics programs

Introduction

The educational system has been identified as a key enabler of Agenda 2030 (Rieckmann, 2017). Adopted by UN member states in 2015, Agenda 2030 aims to address global challenges through 17 Sustainable Development Goals (SDGs).¹ While education is explicitly recognized as a stand-alone Goal (SDG 4), it also plays a critical role in achieving other SDGs

¹ <https://sdgs.un.org/goals>

✉ Mattias Sundemo
mattias.sundemo@economics.gu.se

¹ Department of Economics, University of Gothenburg, Gothenburg, Sweden

² Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, Sweden

³ Asymptotic AI, Lumilogic AB, Gothenburg, Sweden

(Chankseliani & McCowan, 2021). By providing broadened perspectives and developing relevant skills and competencies, education is instrumental in addressing all other global goals, from climate change (SDG 13) and affordable and clean energy (SDG 7) to gender equality (SDG 5) and decent work and economic growth (SDG 8). Subsequently, higher education institutions around the world increasingly recognize the importance of including environmental and societal dimensions in education. Several acclaimed institutions have embedded such ambitions in their mission statements and strategies, for example, Harvard University², “...learning how they [students] can best serve the world,” and Cornell University³, “...applying the knowledge we create for the benefit of society.” Since 2019, Times Higher Education provides an Impact ranking where over 2000 universities are ranked on their contributions to the Sustainable Development Goals⁴.

As ambitions to integrate sustainability into higher education have grown over time, assessing the effects of changes to curricula, as well as other educational activities, has become increasingly important. This assessment is crucial not only for ensuring that developments align with institutional strategies, but also for meeting the growing sustainability reporting requirements of higher education authorities and accreditation bodies (the latter perhaps especially important for Business Schools). While an extensive literature addresses how sustainability is integrated into higher education (for an overview, see Fia et al., 2023), there is generally less guidance on methods for measuring the outcomes of these efforts or, more broadly, on how to measure higher education’s contribution to global common goods (Tian et al., 2024). One important reason for this shortfall is that relevant outcome data are often embedded in text documents, such as curricula, exams questions, exam responses, reports, and, notably, student project papers and Bachelor’s/Master’s theses, which complicates systematic data extraction.

At the same time, technological advances in data-driven text analysis, such as natural language processing (NLP), machine learning, and semantic analysis, have created opportunities for extracting deeper and more comprehensive insights from written text. This has resulted in an increasing number of researchers applying these methods to study research questions which were previously difficult to analyze. For example, there are a number of studies that use such methods to study the integration of sustainability within educational curricula and specifically the alignment with the SDGs (see, e.g., Adams et al., 2023; Lemarchand et al., 2022). However, less attention has been paid to other crucial educational outcome measures, such as how students incorporate sustainability perspectives into their written assignments and theses and how these perspectives have evolved over time, potentially influenced by changes in the curriculum.

In this paper, we develop an “inference engine,” a tool based on advanced keyword-based and data-driven text analysis methods, to assess sustainability content in text. Importantly, the inference engine is based on a methodology that provides a good proxy of educational outcomes by evaluating students’ work from a sustainability perspective, as assessed through alignment with the SDGs. While previous studies have applied similar text analysis methods (primarily NLP) to examine research papers and institutional research priorities (see, e.g., Montenegro De Lima et al., 2020; Abad-Segura and González-Zamar, 2021; Fauzi et al., 2023), our approach shifts the focus from research output to educational output. By analyzing

² <https://college.harvard.edu/about/mission-vision-history> Harvard Mission and Vision.

³ <https://statements.cornell.edu/2019/20190828-community-welcome.cfm> Cornell University Community Statement.

⁴ <https://www.timeshighereducation.com/impactrankings>

student theses, we enable a novel perspective on how the integration of sustainability into education is reflected in student work.

However, data-driven text analysis methods require careful system design and execution. While higher education institutions can, in theory, process and extract data from text sources, effectively leveraging these rapidly evolving methods requires competencies in both data-driven text analysis and sustainability. This combination is essential not only for conducting the analysis but also for accurately interpreting the results. With this in mind, we have aimed to design an inference engine that is transparent and adaptable and produces results that are easy to interpret.

As a case study in this paper, we use the inference engine to assess and compare SDG content in student-written texts (23,758 theses) from the five largest Swedish business schools and universities offering programs in business, economics, and finance. The data extracted using SustainSight provide insights into how sustainability is reflected in students' theses, with a focus on variation across gender, education level (bachelor/master), discipline, and university. Building on recent advances in the causal inference literature (Callaway & Sant'Anna, 2021), we also use a subset of the data to evaluate the causal impact of curriculum reforms, specifically the integration of sustainability into degree requirements, on the likelihood that students incorporate sustainability perspectives into their theses. Finally, we conduct a topic analysis to explore in greater depth what students write about when addressing sustainability-related issues. Our paper contributes methodologically to the assessment of SDG content in student-written texts and empirically to understanding the effects of integrating sustainability into business, economics, and finance education.

The paper is organized as follows. In "[Background and related literature](#)," we introduce the background and review existing literature that analyzes SDGs using written text. We then discuss the specific NLP methods used to map SDGs. In "[How we use NLP to assess sustainability content in theses](#)," we introduce and describe our SDG mapping and inference engine, SustainSight. Following this, in "[Case study: five Swedish universities](#)," we present a case study applying SustainSight to analyze the five largest Swedish business schools and universities that offer degrees in business, economics, and finance. "[Concluding remarks](#)" concludes.

Background and related literature

Sustainability assessment from texts

Any evaluation or assessment of the impact of sustainability integration in educational programs requires a definition of what sustainability as a concept entails. Defining sustainability, however, is quite complex, as can be seen by the extensive body of literature discussing this topic (see, e.g., Kuhlman and Farrington, 2010; Giovannoni and Fabietti, 2013; Farley and Smith, 2020). Hence, assessing whether a body of text incorporates a sustainability perspective requires some degree of judgment. The concept has also evolved from the famous broad Brundtland definition (Brundtland et al., 1987) where *Sustainable Development* was defined as "meeting the needs of the present without compromising the ability of future generations to meet their own needs." Since the adoption of the 17 SDGs, this framework has largely replaced the Brundtland definition in public discourse, particularly within the corporate sphere. While the SDGs are more detailed and address a broader range of challenges than the Millennium Development Goals they replaced, they remain essentially a global "to-do

list” and have been criticized for lacking a strong theoretical foundation and being overly ambitious, and their broad scope hinders focused, effective efforts (Swain, 2018). Additionally, there is an inherent contradiction between the socioeconomic goals (e.g., SDG 8) and the environmental goals (like SDG 13 or 15), particularly given the somewhat vague language used. Despite their shortcomings, the SDGs remain the current framework and roadmap for global efforts on Sustainable Development. They represent a historical consensus, signed by all 193 UN member states in 2015. For the remainder of the paper, we assume that the SDGs capture all relevant aspects of sustainability and will therefore use the terms interchangeably.

There are essentially two ways to assess the sustainability content of written text. One method is to rely on human judgement to infer sustainability relevance from the text which is often time-consuming, and their results can be subjective, as they depend on personal perceptions and experiences. This type of “self-reporting” is commonly used in corporate sustainability reporting, as well as in university accreditation processes and rankings.

The other method for extracting outcome data from text is through automatic analysis, which incorporates both quantitative and qualitative aspects. Quantitative analysis typically determines whether a body of text is relevant to the SDGs in order to aggregate relevant data, while qualitative analysis focuses on extracting topics and context to gain deeper insights. Together, these assessments aim to provide as complete an understanding as possible of how an organization aligns with SDG-related goals and objectives. In the following subsection, we review literature that utilizes automatic assessment analysis to extract information about SDGs in higher education.

SDG mapping in higher education

NLP has been widely used to evaluate sustainability efforts, utilizing the SDGs as a framework. This process, known as *SDG mapping*, typically involves analyzing large text corpora, such as academic papers or policy documents and assigning relevant SDGs to specific segments of text. By determining the relevance of each text to SDGs, researchers attempt to measure how well an institution’s activities align with these goals. The aggregated relevance of all analyzed texts is then used to determine the extent of attention/effort an institution devotes to each SDG. This approach can, if carried out well, enable a systematic and quantifiable assessment of an institution’s work related to sustainability, highlighting strengths and potential gaps.

Systematic and automated approaches in the literature are primarily found in studies that analyze universities’ research output (Smith et al., 2023; Amoros Molina et al., 2023). These typically rely on text analysis of research papers and their bibliometrics (Montenegro De Lima et al., 2020; Prieto-Jiménez et al., 2021; Fauzi et al., 2023), which represent the main scholarly contributions of faculty and researchers. Titles and abstracts are often mapped to relevant SDGs, enabling quantification of SDG-related research as well as insights into institutional sustainability efforts. Some studies further use bibliometric and topic analysis to examine trends and impact of SDG research (Montenegro De Lima et al., 2020; Abad-Segura & González-Zamar, 2021; Fauzi et al., 2023).

An important gap in existing research concerns the types of outputs studied. While research papers are central to universities, they do not necessarily reflect educational outcomes, as they are authored by researchers whose work is guided by their academic interests rather than student learning. Moreover, as researchers secure more funding, they often teach less, limiting students’ interaction with these researchers. This study addresses this key research gap by using student theses as the corpus to assess how sustainability is integrated into higher

education. If universities' sustainability efforts are affecting educational content, such effects can reasonably be expected to appear in these texts.

NLP techniques for SDG mapping

In the literature, two primary techniques are used to assess the SDG relevance of a text segment: keyword-based and data-driven techniques.

Keyword-based techniques

Keyword-based approaches map text to SDGs using predefined keyword lists. For example, Wang et al. (2023) use text mining to construct a contextualized keyword list for identifying SDG-related research, while Körfggen et al. (2018) develop a list of 1000 terms to analyze SDG focus areas in Austrian universities. The widely adopted keyword list by Jayabalasingham et al. (2019) has informed several studies (European Commission. Joint Research Centre, 2023).

Constructing effective keyword lists requires domain expertise, and updates are needed as language evolves. Queries based on these lists allow researchers to analyze texts for SDG alignment. However, poorly designed queries risk false positives (incorrectly identifying irrelevant texts as relevant) or false negatives (failing to identify relevant texts) depending on keyword specificity. For example, a broad term like “energy” may result in a false positive by tagging a thesis on employee engagement as relevant to SDG 7 (Affordable and Clean Energy), while an overly narrow phrase such as “photovoltaic cell efficiency” may cause a false negative by missing broader discussions on renewable energy.

Data-driven techniques

Data-driven techniques use machine learning, deep learning, or generative models to classify text by SDG relevance. These models typically rely on labeled data, such as the annotated corpora from the Open Source SDG (OSDG) initiative (Pukelis et al., 2022, 2020), which include SDG labels and confidence scores.

Machine learning models typically refer to traditional data-driven techniques (Hajikhani & Suominen, 2021, 2022; Morales-Hernandez et al., 2022; Wulff et al., 2023). In contrast, deep learning approaches, which are part of the broader field of machine learning, rely on transformer-based neural network architectures and include models such as BERT (Devlin et al., 2018) and GPT (Han et al., 2021), which are capable of capturing more complex semantic patterns in text. For example, Guisiano et al. (2022) apply deep learning to map text to SDGs.

While deep learning models can capture linguistic nuance, they require substantial domain-specific training data to perform well. Several studies use the OSDG dataset to reduce manual annotation requirements (Angin et al., 2022; Yao et al., 2024). Yu et al. (2024), for example, improve climate-related (SDG 13) classification by using tailored annotations. Other studies generate training data automatically to avoid large-scale annotation (Sovrano et al., 2020; Guisiano & Chiky, 2021).

Importantly, generative models like ChatGPT possess broad reasoning capabilities but may underperform on specific SDG tasks due to their general-purpose design (Raman et al., 2024).

Comparison between keyword-based and data-driven techniques

Keyword-based methods are transparent, require no annotated training data, and can be easily updated to reflect changes in terminology. Their interpretability makes it clear why




Fig. 1 Graphical illustration of the steps in the inference engine SustainSight as applied to analyzing student theses

a text is classified under a given SDG. However, results can vary significantly depending on the chosen keyword set (Armitage et al., 2020), underscoring the importance of careful design. Data-driven techniques, on the other hand, can capture more nuanced patterns beyond predefined terms and adapt to linguistic variation. Yet, they require substantial annotated data and significant computational resources and are sensitive to domain shift (Zhou et al., 2022; Liang et al., 2023). Their black-box nature also limits interpretability.

Hence, while a keyword search approach may be less effective in informal settings, such as social media text mining, where language is more varied and casual, keyword search is well suited for structured texts such as scientific papers and student theses, where language is relatively consistent and precise. Its transparency and rule-based logic make it a reliable baseline for systematic literature reviews and academic database searches (e.g., Scopus, Springer), even alongside large language models (LLM:s).

Finally, the length of input text segments, from words to full documents, is another critical design parameter that affects both approaches (further discussed in the Supplementary Material, Section S1).

How we use NLP to assess sustainability content in theses

In this section, we describe the NLP engine we developed to assess sustainability-related content in theses (henceforth referred to as “SustainSight”). SustainSight⁵ processes a collection of student theses in three steps, as outlined in Fig. 1 and described below.

Step 1: data collection

In the first step, we collect thesis data either from publicly available digital repositories or, when not available, directly from each university. Detailed information about the data collection process for our case study of five Swedish universities is provided in “[Case study: five Swedish universities.](#)”

⁵ Our analysis is based on SustainSight version 1.0.0.

Step 2: SDG classification

SustainSight uses a keyword search mechanism to support transparency and interpretability, as previously discussed.⁶ Drawing on our expertise in sustainability and business education, we adapted keywords from the Elsevier SDG classification framework (Jayabalasingham et al., 2019) to better reflect the language and topics commonly found in theses from business, economics, and finance programs. Each SDG is linked to a tailored keyword list, applied to titles, abstracts, and keyword fields using logical operators (AND, OR, NOT). For example, to identify content on social discrimination under SDG 16, we use the query “discrimination (NOT price)” to exclude references to price discrimination, a common economics term unrelated to discrimination in society. Likewise, “microfinanc* AND poor” is used to classify content under SDG 1 (No Poverty). A detailed description of the process to develop a set of keywords and queries that better match the domain of business, economics and finance education are provided in Supplementary Materials (Section S3).

To assess the robustness of SustainSight, we evaluate its keyword-based method along three dimensions: keyword design, sources of misclassification, and classification performance. Each dimension is discussed thoroughly in Supplementary Material (Section S3). In short, we evaluate the model using blind annotations, meaning that human experts classify abstracts, titles, and user-defined keywords as either related or not related to a specific SDG (without knowing the model’s classification). These human annotations are considered the “ground truth” against which the model is evaluated. The model then classifies the same text segments. The model predicts whether a thesis has an SDG perspective by identifying relevant keywords. Comparing the model’s output to the “ground truth,” we calculate four evaluation metrics: (1) recall, which measures how well relevant SDGs are retrieved; (2) precision, indicating how many retrieved SDGs are actually relevant; (3) accuracy, reflecting overall prediction correctness; and (4) F1 score, which balances precision and recall as their harmonic mean. As further detailed in the Supplementary Material (Section S3), our approach outperforms a state-of-the-art large language model (GPT-4.1) in overall performance, achieving higher recall (with comparable precision). In Section S3, we also discuss the nuances and judgment required, by both human experts and the model, when deciding whether a thesis should be considered to have an SDG perspective. Furthermore, we test our model on a sustainability-focused Master’s program and generally identify 95–100% of the theses as SDG-related (see Supplementary Material, Section S5), which provides further evidence that the model performs well.

In this step, we also extract theses metadata, specifically the authors’ gender,⁷ the level of study (bachelor or master), and disciplinary background (business, economics, or finance).

Step 3: insights

Using the information compiled in the previous steps, we proceed to generate insights from the data. In particular, we produce descriptive statistics, regression outputs, visualize trends, and conduct topic analysis. This type of data can also be combined with other sources for more in-depth analyses, such as causal evaluations of the effect of curricular interventions. These steps are further detailed in “[Case study: five Swedish universities](#).”

⁶ Given the complexity of sustainability, transparent methods are essential. The 16 SDGs (excluding Goal 17, which concerns implementation) provide a widely accepted and comprehensive framework with quantifiable targets. Our approach implicitly assumes these goals capture all aspects of sustainability relevant to higher education.

⁷ The genders of the authors are estimated based on how frequently these names are associated with males and female, specifically focusing on the distinction between male and female students.

Case study: five Swedish universities

The business, economics, and finance disciplines have been criticized for promoting self-interested thinking and lacking ethical foundations (Ghoshal, 2005; Frank et al., 1993; Miragaya-Casillas et al., 2023; Acemoglu et al., 2022; Sundemo & Löfgren, 2024). In response, many business schools have adopted initiatives such as the UN's Principles for Responsible Management Education (PRME), which promotes the integration of ethics, responsibility, and sustainability into curricula, and there is some evidence that universities and business schools increasingly embed SDGs and ethics in business and economics curricula (Weybrecht, 2022). However, we know little about the effects of such reforms. Greater educational exposure to sustainability does not guarantee that students will incorporate sustainability perspectives when given the freedom to choose. In fact, if sustainability interventions are perceived as forced, irrelevant, or of low quality, they may even trigger a backlash.

To shed light on this issue and provide further insights on general trends, this section uses data extracted by SustainSight to assess the sustainability content of Bachelor's and Master's theses from the five largest Swedish universities offering education in business, economics, and finance. Specifically, we examine how sustainability is reflected in student theses across gender, study level, discipline, and institution. We then conduct a causal analysis of how changes to degree requirements with respect to sustainability affect the likelihood that students include SDG perspectives in their theses. Finally, we perform a topic analysis to better understand the themes students address when writing about sustainability-related issues.

Descriptive statistics and initial insights

In sum, our data consist of all digitally available bachelor's and master's theses from students in business, economics, and finance at the University of Gothenburg, Lund University, Stockholm School of Economics, Stockholm University, and Uppsala University. The data cover the period from 2008 to 2022 and include 23,758 theses (a detailed description of the data collection process is provided in the Supplementary Material, Section S2). Each thesis constitutes a unique identifier in this time series and is only observed once.

Table 1 presents descriptive statistics, specifically the number of theses and the share of theses related to at least one SDG, by university, as well as by university and the level of study, disciplinary background, and gender composition of the author team. It is important to note that we only have theses from Stockholm University (SU) in the discipline of business administration, meaning we lack data on economics and finance. Additionally, Uppsala University (UU) does not offer finance. We refer to the Supplementary Material, Section S4, for graphical illustrations of the data, also including more details about the most common SDGs addressed by students and trends over time.

To explore factors associated with the presence of an SDG perspective in theses, we can compare the inclusion of sustainability themes across degree levels (bachelor vs. master), gender composition of the authors, and disciplinary focus. We observe notable differences across gender and disciplinary fields. Across all universities, theses authored exclusively by female students are more likely to engage with SDG perspectives compared to those authored by mixed-gender or all-male groups. Looking at the entire sample, 32% of female-only authored theses relate to at least one SDG, compared to 21.5% of male-only authored theses. This suggests that female students are nearly 1.5 times more likely than male students

Table 1 Descriptive statistics, number of theses, and share of theses relating to at least one SDG (2008–2022)

	University of Gothenburg	Stockholm School of Econ.	Lund University	Stockholm University	Uppsala University	Total
All theses	5034 (30%)	4067 (23%)	5755 (28%)	3582 (23%)	5320 (25%)	23,758
<i>By gender</i>						
Male only	1859 (24%)	1699 (18%)	2187 (25%)	1055 (18%)	2129 (19%)	8929
Female only	1752 (35%)	1196 (28%)	1713 (37%)	1317 (26%)	1751 (31%)	7729
Mixed	782 (32%)	584 (25%)	1025 (20%)	630 (24%)	753 (21%)	3774
Unknown	641 (32%)	588 (23%)	830 (23%)	580 (25%)	687 (28%)	3326
<i>By degree level</i>						
Bachelor	2863 (33%)	1799 (22%)	2841 (28%)	2006 (24%)	3421 (23%)	12,930
Master	2171 (27%)	2268 (23%)	2914 (27%)	1576 (22%)	1899 (28%)	10,828
<i>By discipline</i>						
Business	4002 (27%)	1584 (22%)	3235 (19%)	3582 (23%)	4148 (19%)	16,551
Economics	657 (54%)	666 (49%)	2300 (41%)	–	1172 (45%)	4795
Finance	375 (23%)	1817 (14%)	220 (20%)	–	–	2412

Note: The numbers provided are total thesis counts within the given category. Percentages in parentheses report the proportion of theses that our model identified as related to *at least one* SDG. “Mixed” denotes teams containing both male and female authors, whereas “Unknown” refers to teams for which at least one author’s gender cannot be identified. An en dash symbol (–) indicates the absence of observations in that cell

to include an SDG perspective in their theses. Similarly, economics theses are substantially more likely to include sustainability content than those from business or finance programs. In contrast, degree level (bachelor vs. master) does not appear to be systematically associated with SDG inclusion.

These patterns are further explored in Table 2, where we report regression estimates summarizing these associations. We estimate a logistic regression model (reporting marginal effects) for the binary outcome of whether a thesis includes at least one SDG perspective, and an OLS model where the dependent variable is the total count of SDGs referenced, based on SustainSight’s classification. The independent variables include gender composition, degree level, discipline, university affiliation, and time periods. While the underlying data allow us to examine associations, they do not permit causal inference. Nonetheless, the regression results are consistent with the descriptive findings.

Specifically, theses authored by female-only groups are associated with a 5 percentage point higher likelihood of including an SDG perspective than those authored by mixed-gender pairs, while male-only theses are associated with a 7 percentage point lower likelihood. This implies that a thesis written by females-only are almost 12 percentage points more likely to include an SDG perspective than a thesis written by males-only. Disciplinary differences are particularly pronounced: finance theses are 29 percentage points less likely, and business theses 23 percentage points less likely, to include an SDG perspective compared to economics theses.

These associations suggest that gender and disciplinary background are strongly correlated with the inclusion of sustainability perspectives, an observation that aligns with previous findings in the literature. While earlier evidence was somewhat inconclusive, more recent studies suggest that women tend to report stronger environmental attitudes, concerns, and behaviors than men (see, e.g., Gifford & Nilsson, 2014; Ramstetter & Habersack, 2020; Gyurián Nagy,

Table 2 Factors associated with SDG engagement in theses

	At least one SDG		Total SDG	
	Coef	SE	Coef	SE
<i>University (Stockholm University as reference)</i>				
Lund University	−0.052***	0.010	−0.101***	0.018
Uppsala University	−0.038***	0.009	−0.067***	0.018
Stockholm School of Economics	−0.022*	0.011	−0.051**	0.021
University of Gothenburg	0.033***	0.009	0.089***	0.018
<i>Gender (females and males writing together as reference)</i>				
Females only	0.047***	0.008	0.099***	0.016
Males only	−0.070***	0.009	−0.133***	0.016
Unknown gender	0.005	0.010	0.020	0.020
<i>Discipline (Econ as reference)</i>				
Business	−0.232***	0.006	−0.438***	0.014
Finance	−0.288***	0.012	−0.520***	0.023
<i>Bachelor vs Master (Master as reference)</i>				
Bachelor	−0.007	0.006	−0.011	0.011
<i>Time (2008–2012 as reference)</i>				
2013–2017	0.027***	0.007	0.046***	0.013
2018–2022	0.092***	0.007	0.185***	0.013
Observations	23,758		23,758	

Note: This summarizes results from two regressions with all listed co-variables and time dummies. The first model (“At least one SDG”) is a logistic regression; reported coefficients are average marginal effects. The second model (“Total SDG”) is an OLS regression estimating how the number of SDGs mentioned in a thesis correlates with the same set of covariates

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

2025). Regarding disciplinary differences, there is an extensive body of literature comparing pro-sociality (vs. self-interest) and environmental attitudes across student groups, particularly focusing on business and economics students compared to other disciplines (Lang, 2011; Hellmich, 2019; Miragaya-Casillas et al., 2023; Sundemo & Löfgren, 2024). A key takeaway from this literature is that business and economics majors tend to exhibit lower levels of environmental concern and pro-sociality than students in other fields. Moreover, Sundemo and Löfgren (2024) provide suggestive evidence of differences even within related disciplines, indicating that economics students demonstrate stronger prosocial values than their peers in business and finance programs. While our study does not aim to explain the underlying mechanisms driving these differences, such as self-selection, educational impact, or socialization, our results are consistent with these prior findings.

We also observe statistically significant differences across universities. Degree level, on the other hand, appears to play a limited role. However, the fact that university affiliation is also significantly associated with SDG inclusion raises further questions. One possible explanation is that universities differ in their integration of sustainability into their curricula. Further, inclusion of SDG perspectives seems to increase over time. It is 9 percentage points more likely that a thesis written in the period 2018–2022 include an SDG perspective as compared to the first five year period. In the next section, we build on these observations to

explore whether we can draw more robust inferences about the potential effect of curriculum-level sustainability integration on student thesis content.

Finally, it can be noted that the OLS count model yields qualitatively similar results to the logistic regression, reinforcing the observed patterns.

Estimating the causal effect of exposure on outcomes

This section will explore the following question: has the introduction of degree requirements for sustainability in traditional Swedish business and economics programs increased the likelihood that students choose to relate their theses to topics covered by the SDGs? If business, economics, and finance students increasingly engage with these topics, it may contribute to strengthening society's capacity to address sustainability challenges.

Data and treatment definition

In “*Descriptive statistics and initial insights*,” we included the full sample in the analysis, without controlling for exposure to sustainability content in the education. Some students (albeit few) enroll in sustainability dedicated programs, some pursue a degree in business, economics, or finance while enrolled in a program from a different field (eg. engineering or political science), and yet others (a large proportion of students), put together a degree using freestanding courses that they choose themselves, often from different fields.

In this section, we will restrict the sample to traditional business and economics bachelor programs, for which each of the five universities has one. Between 2013 and 2019, all programs introduced degree requirements for sustainability (see Table 3). We argue that students enrolling in these traditional business and economics programs are comparable and constitute a suitable control group for one another.

In order to enable this analysis, we matched thesis titles to admission data from the database Ladok.⁸ By doing so, we can control for whether or not sustainability were integrated in program degree requirements and test if increased sustainability exposure predict increases in the share of theses that relate to the topics covered by the SDGs. We consider a thesis (program) as treated if at least one degree requirement explicitly related to sustainability when the student enrolled in the program. Hence, we observe the first treated cohorts 3 years after the new requirements were implemented, that is, when the first cohort writes their theses.

For the matched dataset, where each thesis is paired with a traditional business and economics program, we have 6442 theses from 2011 to 2022. This includes all theses we could link from our original collection to the Ladok records supplied by each university. Since Swedish universities typically require group work (up to four authors per thesis, with three being common), the total number of students exceeds the thesis count (as in the full dataset). Annual thesis counts are shown in Table 4.

⁸ The database, Ladok, is Sweden's nationwide student information system, used by universities and colleges to record enrollments, course registrations, grades and other academic results. A considerable effort, including manual work, was required to ensure proper matching between a theses title and an educational program as thesis titles in the different systems not always perfectly matched Ladok titles, sometimes the full title was cut short in one of the databases for example. More details about this can be found in Supplementary Material, Section S2.2.

Table 3 Sustainability interventions in Swedish bachelor programmes

Institution (programme)	Introduced	First cohort treated	Type of intervention
University of Gothenburg <i>Bachelor's Programme in Business and Economics</i>	2013	2016	Degree requirement (program learning goals) and 3 dedicated Sustainability Days ¹
Stockholm University <i>Bachelor's Programme in Business Administration</i>	2017	2020	Degree requirement (program learning goals) and “theme” semesters (sustainability and ethics) ²
Lund University <i>Bachelor's Programme in Business Administration and Economics</i>	2017	2020	General degree requirement (program learning goals referencing overall university strategy for sustainability) ³
Uppsala University <i>Bachelor's Programme in Business and Economics</i>	2019	2022	Degree requirement (program learning goal) ⁴
Stockholm School of Economics <i>Bachelor's Programme in Business and Economics</i>	2016	2019	Introduction of mandatory “Global Challenges module” parallel to traditional studies ⁵

Note: ¹Three learning outcomes focusing on environmental and social aspects. In addition, yearly sustainability days were introduced with themes related to sustainability challenges, responsibilities, and solutions. ²Two “theme” semesters: one on sustainability, one on ethics, each integrated into the core curriculum, complemented with two sustainability courses. ³Curriculum complemented with a learning outcome that reference the overall university sustainability strategy. ⁴A single program-level learning goal on sustainability was added to the official degree requirements in 2019, making all programs “treated” in 2022. ⁵A so-called Global Challenges module was introduced, running alongside standard courses; the first part (6 HEC) introduces topics such as global inequality and corporate responsibility

Table 4 Number of bachelor and master theses by year and university (2011–2022)

Year	U Gothenburg	Stockholm Univ.	Lund Univ.	Uppsala Univ.	SSE
2011	68	80	103	86	86
2012	68	74	125	30	113
2013	86	70	180	108	88
2014	89	74	183	175	101
2015	91	79	222	159	90
2016	81	88	159	131	90
2017	103	101	135	139	92
2018	103	89	151	126	97
2019	76	54	133	111	105
2020	101	104	115	118	105
2021	123	145	144	106	72
2022	183	68	50	109	107
Total	1 172	1 026	1 700	1 398	1 146

Note: SSE Stockholm School of Economics. Counts include our original dataset matched to Ladok records for traditional business and economics programs described in Table 3

Empirical strategy

To estimate the causal effect of exposure to sustainability content on our outcome variable, we exploit the fact that the five traditional bachelor programs in business and economics adopted sustainability degree requirements in different calendar years (Table 3). This lets us apply a so-called *staggered difference-in-differences* (Goodman-Bacon, 2021). The foundational idea of any DiD is that a group cannot simply be compared to itself before and after treatment. One needs a suitable control group to estimate a plausible counterfactual (what would have happened in the absence of treatment). In our case, one might think that independent of inclusion of degree requirement for sustainability, students will over time be more likely relate to topics covered by the SDGs (regressions on the whole sample do find a positive association with time as shown in Table 2). To understand if such trends are due to interventions in education, or rather an overall trend in society, we employ a staggered difference-in-difference model where we in each calendar year compare treated schools (that have implemented a degree requirement) with schools that have not. We describe this formally below.

Let g denote the first cohort treated in school s and t denote calendar year. The cohort- and time-specific average treatment effect on the treated (ATT) is then as follows:

$$ATT(g, t) = \mathbb{E}[Y_{st}(1) - Y_{st}(0) \mid G_s = g, t], \quad (1)$$

where $Y_{st}(1)$ is the SDG-thesis share under treatment and $Y_{st}(0)$ is the counterfactual without treatment. Naturally, a school (program) cannot be both treated and untreated, and we therefore approximate $Y_{st}(0)$ with outcomes from schools that are not-yet treated in the same year. We then aggregate the $ATT(g, t)$ across cohorts to get the overall effect.

Pre-treatment trends and homogeneous treatment effects

A staggered DiD, traditionally estimated with a so-called two-way fixed effects (TWFE) regression, relies on the same core assumption as the simple two-period design, namely that individuals that become treated in a given calendar year were not on a different trajectory than individuals that stay untreated in that same year. However, recent work by Goodman-Bacon (2021); Sun and Abraham (2021) point out that the TWFE rely on another important assumption, namely that treatment effects are constant across cohorts and time. This is because the TWFE uses early treated units as control for later treated units. Therefore, if the treatment effect grows over time, as is the case in many settings, this approach will generate biased estimates. In our case, there is little reason to assume that treatment effects are strictly homogeneous across cohorts and over time. We therefore make use of the Callaway and Sant'Anna staggered difference-in-differences estimator (Callaway & Sant'Anna, 2021). This method compares each treated cohort only with programs that are still untreated in the same calendar year and then aggregates the cohort-specific effects to an average treatment effect on the treated.

Event-study specification

We estimate the causal effect of introducing sustainability-related degree requirements on the probability that a thesis relates to the SDGs, using the event study approach outlined below, which accounts for heterogeneous treatment effects (Callaway & Sant'Anna, 2021).

For thesis i written in school s and year t , we estimate

$$Y_{ist} = \alpha + \sum_{k \neq -1} \beta_k \mathbf{1}\{\text{event_time}_{st} = k\} + \lambda_s + \gamma_t + \varepsilon_{ist}, \quad (2)$$

with

$$\text{event_time}_{st} = \text{year}_t - \text{treat_year}_s,$$

That is, $k = 0$ is the first treated cohort, and $k = -1$ serves as the omitted baseline (1 year before treatment). School fixed effects (λ_s) account for time-invariant differences across schools, while calendar-year fixed effects (γ_t) control for common shocks and time trends. The coefficients β_k estimate the dynamic impact of exposure on outcome, the share of SDG-related theses. We cluster standard errors at the program level as treatment is assigned at that level.

Pre-treatment trends

As noted above, any DiD approach relies on the parallel trends assumption: that, in the absence of the policy, treated programs would have followed the same trajectory as those that remained untreated. Figure 2 therefore plots, by calendar year, the share of SDG-related theses only for the *untreated programs* in each year. As can be seen, there is no discernible upward or downward trend, and the average hovers around 20% from 2011 to 2021, supporting the parallel trends assumption.

Results

Figure 3 plots the raw average share of theses relating to at least one SDG relative to the time of treatment. As can be seen, the period up to treatment, while a little noisy, does not indicate




Fig. 2 Share of theses relating to at least one SDG over time for the not-yet-treated sample




Fig. 3 Raw shares of theses relating to at least one SDG relative to year of treatment

any particular trend. Post treatment, however, the share of theses relating to at least one SDG increases from roughly 20 to about 30% after 3 years and then continues to rise further to over 40%.

While Fig. 3 is informative, it does not identify the causal effect because two confounding factors remain: time effects and school effects. We therefore proceed to present the results from the event study specification (Eq. 2), which controls for program-specific effects, accounts for year-specific shocks, and is robust to heterogeneous treatment effects (e.g., effects that grow over time). The corresponding estimates are shown in Fig. 4. Using the Callaway and Sant’Anna estimator (Callaway & Sant’Anna, 2021), we find that 2 years after treatment, the probability that a thesis relates to at least one SDG increases by approximately 9 percentage points, rising to 16 percentage points after 3 years.⁹ This percentage point change constitutes a substantial increase: relative to the pooled pre-treatment mean of roughly 20%, the effect 2 years post-treatment corresponds to an increase of about 45% and 80% 3 years after treatment.

Addressing potential threats to causal inference

Because treatment is assigned at the program level, statistical inference must account for the possibility that error terms are correlated within universities across years. We therefore cluster standard errors at school (program) level. However, with only five clusters, conventional cluster-robust standard errors can be significantly downward-biased (Cameron & Miller, 2015). We employ the recommended approach and use wild cluster bootstrap standard errors. Still, while point estimates are unaffected, standard errors should be interpreted with caution.

⁹ Beyond 3 years post-treatment, only one treated program (GU) and one not-yet-treated program (Uppsala) remain, making it infeasible to compute standard errors.




Fig. 4 Staggered difference-in-differences: probability of a theses relating to at least one SDG. *Note:* This displays the average treatment effects on the treated (ATT), estimated by a staggered difference-in-differences (Callaway–Sant’Anna) event study on the binary indicator “relates to at least one SDG.” Point estimates are relative to the year before treatment and the vertical bars display 95% confidence intervals. Each cohort is compared only to not-yet-treated programs, and effects are aggregated across cohorts

A more fundamental concern is whether the observed effect truly reflects increased curricular exposure or if it arises from selection. Perhaps students select into (or out of) treated programs based on preferences for a more sustainability-focused curriculum? While possible, this seems unlikely, given that curriculum updates are typically not salient to students during the application process. Further, students with a strong interest in sustainability likely choose one of the sustainability-focused programs offered at these business schools.¹⁰ Still, we test for selection along two observable dimensions. Table 2 shows that (i) theses written by female-only students are more likely to be SDG-related and (ii) economics theses have a higher SDG share than business or finance theses. If selection were driving the results, treated programs should likely attract more women and/or see a rising share of economics theses after adoption. A difference-in-differences regression with program and year find no support for this: the change in the female share is virtually zero ($p = 0.97$), and the share of economics theses, if anything, falls slightly but insignificantly ($p = 0.60$). Because we find no evidence of selection on relevant observables, we conclude that increased exposure is the more plausible explanation for the positive effect we observe.

Why the lagged effect?

There are several reasons to expect a gradual and increasing treatment effect over time. First, cohort lag is likely to affect the outcome: not all individuals from the first treated cohort will

¹⁰ University of Gothenburg, Stockholm University, Lund University, and Uppsala University all have different versions of such dedicated sustainability programs.

graduate on time, and these individuals will therefore be observed at a later stage than 3 years after enrollment. Likewise, outcomes observed for the first treated cohort may include work by students from earlier cohorts who did not graduate on time and were enrolled under the previous curriculum without sustainability degree requirements. Second, the introduction of new degree requirements can be seen as the beginning of a broader development process. Early cohorts may experience both lower quantity and quality of implementation, while program content and course modules are likely to improve as teachers and supervisors gain experience and knowledge.

Data-driven topic analysis of sustainability theses

To gain a deeper understanding of the recurring topics and themes that Business School students explore within the realm of sustainability, we conduct a data-driven topic analysis based on clustering of thesis abstracts, title and keywords. In this analysis, we utilize all the theses from the full sample that was identified as related to at least one SDG (see Table 1).¹¹

To achieve this, we run a clustering analysis algorithm to group related theses and summarize their topics. Specifically, we first transform each thesis into a numerical vector representation (Su et al., 2022). We then apply a dimensionality reduction algorithm (McInnes et al., 2018) to project these vectors into a two-dimensional space. We applied Gaussian mixture model (GMM) and HDBSCAN clustering algorithms with various hyperparameters (such as number of clusters for GMM and minimum cluster size for HDBSCAN) and manually evaluated topic quality by reviewing sampled theses. Based on this assessment, we selected GMM with 20 clusters to group similar theses in the two-dimensional space.

Table 5 presents these clusters along with their corresponding topics, the total number of theses in each cluster across all five universities (count), their relative size (share of all theses), and the share of each cluster at the individual universities.

The results show that SDG theses cover a wide range of topics. The five largest clusters, sustainable consumption, global development, innovation, CSR, and sustainability reporting, account for nearly half of all SDG theses. These clusters reflect a mix of qualitative and quantitative approaches across the disciplines of business, economics and finance. The five universities show some variation in topic focus. The largest cluster at GU, SSE, and SU is “Sustainable consumption and consumer attitudes” (cluster 0), while at LU and SU, “Global development, FDI, and trade” (Cluster 1) dominates. At UU, Cluster 1 and the cluster “Innovation and entrepreneurship” is of nearly equal size.

Our underlying hypothesis is that topics do not emerge arbitrarily; instead, each reflects a distinct combination of SDGs, shaped by systematic patterns in the data. In other words, topics are not random mixtures but are meaningfully aligned with certain SDGs more than others. To statistically test this assumption (whether the distribution of SDG associations differs significantly across topics), we conducted a simple chi-square test of independence of topics and SDGs. The results support the hypothesis that certain SDGs are over- or under-represented in specific topics ($p < 0.001$), reinforcing the idea that thematic clusters exhibit distinct sustainability orientations.

The link between specific SDGs and the topics is visualized in Fig. 5. The figure presents a set of radial plots, one for each of the 20 thematic clusters, visualizing the distribution of SDG perspectives within each cluster.¹² Each cone represents the share of theses in the cluster that relate to a specific SDG, grouped into three overarching domains: biosphere (green),

¹¹ We applied the method only to English-language theses due to model limitations.

¹² Each thesis can contribute to multiple SDGs.

Table 5 Topics from the cluster analysis

#	Topic/theme	Count	Share	GU	SSE	LU	SU	UU
0	Sustainable consumption and consumer attitudes	655	13%	14%	11%	9%	23%	10%
1	Global development, FDI, and trade	498	10%	7%	8%	18%	2%	11%
2	Innovation and entrepreneurship	442	9%	8%	10%	7%	10%	11%
3	Corporate social responsibility (CSR)	395	8%	10%	6%	4%	13%	8%
4	Sustainability reporting and disclosure	335	7%	6%	3%	5%	12%	8%
5	Sustainable finance: SRI and green bonds	284	6%	5%	6%	4%	9%	5%
6	Sustainable finance: ESG and financial performance	284	6%	5%	8%	4%	8%	4%
7	Sustainability transition in Sweden: Policy, innovation, and sectoral practices	253	5%	12%	3%	3%	3%	3%
8	Impact of gender dynamics and socioeconomic factors on economic development	224	4%	4%	4%	7%	<1%	6%
9	Healthcare management and health outcomes	212	4%	3%	4%	6%	3%	4%
10	Housing markets dynamics	210	4%	3%	6%	5%	3%	4%
11	Labor market dynamics, education, and migration in a European context	186	4%	2%	3%	6%	0%	7%
12	Macroeconomic outcomes and financial market dynamics	171	3%	2%	4%	6%	3%	2%
13	Circular economy, waste, and recycling	166	3%	6%	3%	2%	2%	3%
14	Employee well-being, health, and organizational change	156	3%	3%	4%	2%	5%	2%
15	Income inequality, immigration, crime, and social trust	154	3%	1%	4%	4%	0%	5%
16	Climate and environmental policy, economic growth, and emissions outcomes	132	3%	2%	4%	4%	0%	2%
17	Gender and equality in business	128	3%	2%	4%	2%	4%	2%
18	Renewable energy markets	94	2%	2%	2%	2%	<1%	2%
19	Sustainability in business and smart cities	58	1%	<1%	1%	1%	<1%	2%
	Total	5037	100%					

GU University of Gothenburg, *SSE* Stockholm School of Economics, *LU* Lund University, *SU* Stockholm University, *UU* Uppsala University




Fig. 5 Topics and SDGs

society (blue), and economy (orange). The length of each cone indicates the percentage of theses that relate to the SDG, allowing for comparisons both within and across clusters. The plots reveal distinct patterns: some clusters, such as “Innovation and entrepreneurship” or “Sustainability in business and smart cities,” are linked to a broad set of SDGs, while others, like “Renewable energy markets,” show a concentrated alignment with specific SDGs (e.g., SDG 7). The variation in SDG profiles suggests that different thematic areas reflect differing

degrees of integration across sustainability dimensions, likely shaped by their underlying problem framings.

Concluding remarks

In this study, we provide an overview of the literature on mapping SDGs in higher education and explain various NLP techniques used for SDG mapping. Further, we present an inference engine based on NLP that we developed to assess sustainability content in student theses.

Analyzing student theses in business, economics, and finance at five Swedish universities suggests that gender and disciplinary background are strongly correlated with the inclusion of sustainability perspectives. Theses written by female-only authors are 50% more likely to include an SDG perspective than those written by male-only authors. Regarding disciplinary differences, finance theses are 29 percentage points less likely, and business theses 23 percentage points less likely, to include an SDG perspective compared to economics theses.

For a subset of the theses, we have information about the authors' educational programs as well as data on the sustainability degree requirements. Using this data and the fact that the five traditional business and economics programs introduced sustainability degree requirements at different points in time, we can conduct a causal analysis (a so-called staggered difference-in-differences). We find that the introduction of a degree requirement focused on sustainability significantly increases the likelihood that students include sustainability perspectives in their theses. We find an increase of 9 percentage points (around 45%) at year +2 and 16 points (around 80%) at year +3, relative to the first treated cohort.

While our analysis focuses on Sweden and we cannot formally test for external validity, it is worth noting that business, economics, and finance education is highly standardized globally, with similar textbooks, theories, and core courses. We therefore consider it likely that our results offer insights applicable to other national contexts.

In conclusion, we argue that applying NLP to student theses offers a promising method for assessing educational outcomes, particularly the integration of sustainability perspectives in higher education. However, when using such methods, it is important to consider and mitigate biases, such as incomplete or uneven keyword coverage, through regular expert reviews (human-in-the-loop).

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10734-025-01526-9>.

Author contribution Mattias Sundemo, Åsa Löfgren, and Yinan Yu contributed to the conceptualization, methodology, data curation, formal analysis, visualization, and writing (original draft, review, and editing). Samuel Scheidegger contributed to visualization and data curation.

Funding Open access funding provided by University of Gothenburg. Financial support from the UGOT Centre for Collective Action Research and the Familjen Kamprads stiftelse, project 20230142, Policies for a Sustainable Society is gratefully acknowledged.

Data Availability Data are available upon request.

Declarations

Ethics approval and consent to participate Not applicable

Conflict of interest The model developed and presented in this paper was created for research purposes, and the findings and analyses reflect its research-phase use. The authors note that there is potential for future commercialization of the model and associated pipeline.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abad-Segura, E., & González-Zamar, M. D. (2021). Sustainable economic development in higher education institutions: A global analysis within the SDGs framework. *Journal of Cleaner Production*, 294, 126133. <https://doi.org/10.1016/j.jclepro.2021.126133>. <https://linkinghub.elsevier.com/retrieve/pii/S095965262100353X>
- Acemoglu, D., He, A., & Le Maire, D. (2022). *Eclipse of rent-sharing: The effects of managers' business education on wages and the labor share in the US and Denmark*. Tech. rep., National Bureau of Economic Research
- Adams, T., Jameel, S. M., & Goggins, J. (2023). Education for sustainable development: Mapping the SDGs to university curricula. *Sustainability*, 15(10), 8340.
- Amoros Molina, A., Hellden, D., Alfven, T., et al. (2023). Integrating the United Nations Sustainable Development Goals into higher education globally: A scoping review. *Global Health Action*, 16(1), 2190649.
- Angin, M., Taşdemir, B., & Yılmaz, C. A., et al. (2022). A RoBERTa approach for automated processing of sustainability reports. *Sustainability*, 14(23), 16139. <https://doi.org/10.3390/su142316139>. <https://www.mdpi.com/2071-1050/14/23/16139>
- Armitage, C. S., Lorenz, M., & Mikki, S. (2020). Mapping scholarly publications related to the Sustainable Development Goals: Do independent bibliometric approaches get the same results? *Quantitative Science Studies*, 1(3), 1092–1108. https://doi.org/10.1162/qss_a_00071. <https://direct.mit.edu/qss/article/1/3/1092-1108/96106>
- Brundtland, G., et al. (1987). Our common future. *The world commission on environment 1 and development* (pp. 45–65).
- Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
- Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2), 317–372.
- Chankseliani, M., & McCowan, T. (2021). Higher education and the sustainable development goals. *Higher Education*, 81(1), 1–8.
- Devlin, J., Chang, M. W., Lee, K., et al. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805)
- European Commission. Joint Research Centre. (2023). Measuring the impact of digital technologies on sustainable development from scientific literature: A natural language processing analysis. Publications Office, LU. <https://data.europa.eu/doi/10.2760/64241>
- Farley, H. M., & Smith, Z. A. (2020). *Sustainability: If it's everything, is it nothing?* Routledge.
- Fauzi, M. A., Abdul Rahman, A. R., & Lee, C. K. (2023). A systematic bibliometric review of the United Nation's SDGs: Which are the most related to higher education institutions? *International Journal of Sustainability in Higher Education*, 24(3), 637–659. <https://doi.org/10.1108/IJSHE-12-2021-0520>. <https://www.emerald.com/insight/content/doi/10.1108/IJSHE-12-2021-0520/full/html>
- Fia, M., Ghasemzadeh, K., & Paletta, A. (2023). How higher education institutions walk their talk on the 2030 agenda: A systematic literature review. *Higher Education Policy*, 36(3), 599–632.
- Frank, R. H., Gilovich, T., & Regan, D. T. (1993). Does studying economics inhibit cooperation. *Journal of Economic Perspectives*, 7(2), 159–171.
- Ghoshal, S. (2005). Bad management theories are destroying good management practices. *Academy of Management Learning & Education*, 4(1), 75–91.

- Gifford, R., & Nilsson, A. (2014). Personal and social factors that influence pro-environmental concern and behaviour: A review. *International Journal of Psychology*, 49(3), 141–157.
- Giovannoni, E., & Fabietti, G. (2013). What is sustainability? A review of the concept and its applications. *Integrated reporting: Concepts and cases that redefine corporate accountability* (pp. 21–40).
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277. <https://doi.org/10.1016/j.jeconom.2021.03.014>. <https://www.sciencedirect.com/science/article/pii/S0304407621001445>, themed Issue: Treatment Effect 1.
- Guisiano, J., & Chiky, R. (2021). Automatic classification of multilabel texts related to Sustainable Development Goals (SDGs). In *TECHENV EGC2021*.
- Guisiano, J. E., Chiky, R., & De Mello, J. (2022). SDG-Meter: A deep learning based tool for automatic text classification of the sustainable development goals. In *Asian conference on intelligent information and database systems*. Springer, pp. 259–271.
- Gyurián Nagy, N. (2025). Gender differences in environmental attitudes: An analysis using the NEP scale. *Gender Issues*, 42(1), 5.
- Hajikhani, A., & Suominen, A. (2021). The interrelation of sustainable development goals in publications and patents: A machine learning approach. *AI+ Informetrics* (pp. 183–193).
- Hajikhani, A., & Suominen, A. (2022). Mapping the Sustainable Development Goals (SDGs) in science, technology and innovation: Application of machine learning in SDG-oriented artefact detection. *Scientometrics*, 127(11), 6661–6693.
- Han, X., Zhang, Z., Ding, N., et al. (2021). Pre-trained models: Past, present and future. *AI Open*, 2, 225–250.
- Hellmich, S. N. (2019). Are people trained in economics “different,” and if so, why? A literature review. *The American Economist*, 64(2), 246–268.
- Jayabalasingham, B., Boverhof, R., Agnew, K., et al. (2019). Identifying research supporting the United Nations Sustainable Development Goals. *Mendeley Data*, 1(1).
- Körfggen, A., Förster, K., Glatz, I., et al. (2018). It’s a hit! Mapping Austrian research contributions to the Sustainable Development Goals. *Sustainability*, 10(9), 3295. <https://doi.org/10.3390/su10093295>. <https://www.mdpi.com/2071-1050/10/9/3295>
- Kuhlman, T., & Farrington, J. (2010). What is sustainability? *Sustainability*, 2(11), 3436–3448.
- Lang, K. B. (2011). The relationship between academic major and environmentalism among college students: Is it mediated by the effects of gender, political ideology and financial security? *The Journal of Environmental Education*, 42(4), 203–215.
- Lemarchand, P., McKeever, M., MacMahon, C., et al. (2022). A computational approach to evaluating curricular alignment to the united nations sustainable development goals. *Frontiers in Sustainability*, 3, 909676. <https://doi.org/10.3389/frsus.2022.909676>. <https://www.frontiersin.org/articles/10.3389/frsus.2022.909676/full>
- Liang, J., He, R., & Tan, T. (2023). A comprehensive survey on test-time adaptation under distribution shifts. [arXiv:2303.15361](https://arxiv.org/abs/2303.15361)
- McInnes, L., Healy, J., & Melville, J. (2018). Umap: Uniform manifold approximation and projection for dimension reduction. [arXiv:1802.03426](https://arxiv.org/abs/1802.03426)
- Miragaya-Casillas, C., Aguayo-Estremera, R., & Ruiz-Villaverde, A. (2023). University students, economics education, and self-interest: A systematic literature review. *International Review of Economics Education*, 43, 100266.
- Montenegro De Lima, C. R., Coelho Soares, T., Andrade De Lima, M., et al. (2020). Sustainability funding in higher education: A literature-based review. *International Journal of Sustainability in Higher Education*, 21(3), 441–464. <https://doi.org/10.1108/IJSHE-07-2019-0229>. <https://www.emerald.com/insight/content/doi/10.1108/IJSHE-07-2019-0229/full/html>
- Morales-Hernandez, R. C., Jaguey, J. G., & Becerra-Alonso, D. (2022). A comparison of multi-label text classification models in research articles labeled with sustainable development goals. *IEEE Access*, 10, 123534–123548. <https://doi.org/10.1109/ACCESS.2022.3223094>. <https://ieeexplore.ieee.org/document/9954368/>
- Prieto-Jiménez, E., López-Catalán, L., López-Catalán, B., et al. (2021). Sustainable development goals and education: A bibliometric mapping analysis. *Sustainability*, 13(4), 2126. <https://doi.org/10.3390/su13042126>. <https://www.mdpi.com/2071-1050/13/4/2126>
- Pukelis, L., Bautista-Puig, N., Statulevičiūtė, G., et al. (2022). Osdg 2.0: A multilingual tool for classifying text data by un sustainable development goals (sdgs). [arXiv:2211.11252](https://arxiv.org/abs/2211.11252)
- Pukelis, L., Puig, N. B., Skrynik, M., et al. (2020). Osdg—open-source approach to classify text data by un sustainable development goals (sdgs). [arXiv:2005.14569](https://arxiv.org/abs/2005.14569)
- Raman, R., Lathabai, H. H., Mandal, S., et al. (2024). ChatGPT: Literate or intelligent about UN sustainable development goals? *PLoS One*, 19(4), Article e0297521. <https://doi.org/10.1371/journal.pone.0297521>

- Ramstetter, L., & Habersack, F. (2020). Do women make a difference? Analysing environmental attitudes and actions of members of the European parliament. *Environmental Politics*, 29(6), 1063–1084. <https://doi.org/10.1080/09644016.2019.1609156>
- Rieckmann, M. (2017). *Education for sustainable development goals: Learning objectives*. UNESCO publishing
- Smith, T. B., Vacca, R., Mantegazza, L., et al. (2023). Discovering new pathways toward integration between health and sustainable development goals with natural language processing and network science. *Globalization and Health*, 19(1), 44. <https://doi.org/10.1186/s12992-023-00943-8>. <https://globalizationandhealth.biomedcentral.com/articles/10.1186/s12992-023-00943-8>
- Sovrano, F., Palmirani, M., & Vitali, F. (2020). Deep learning based multi-label text classification of UNGA resolutions. In *Proceedings of the 13th International Conference on Theory and Practice of Electronic Governance*. ACM, Athens Greece, pp. 686–695. <https://doi.org/10.1145/3428502.3428604>. <https://dl.acm.org/doi/10.1145/3428502.3428604>
- Su, H., Shi, W., Kasai, J., et al. (2022). One embedder, any task: Instruction-finetuned text embeddings. [arXiv:2212.09741](https://arxiv.org/abs/2212.09741)
- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
- Sundemo, M., & Löfgren, Å. (2024). Do business and economics studies erode prosocial values? *Southern Economic Journal*.
- Swain, R. B. (2018). A critical analysis of the sustainable development goals. *Handbook of sustainability science and research* (pp. 341–355)
- Tian, L., Feng, Z. L., & Liu, N. C. (2024). Measuring global common goods in higher education: Dimensions and potential indicators. *Higher Education* (pp. 1–16)
- Wang, W., Kang, W., & Mu, J. (2023). Mapping research to the Sustainable Development Goals (SDGs). <https://doi.org/10.21203/rs.3.rs-2544385/v2>. <https://www.researchsquare.com/article/rs-2544385/v2>
- Weybrecht, G. (2022). Business schools are embracing the sdgs - but is it enough? - How business schools are reporting on their engagement in the sdgs. *The International Journal of Management Education*, 20(1), 100589. <https://doi.org/10.1016/j.ijme.2021.100589>. <https://www.sciencedirect.com/science/article/pii/S1472811721001385>
- Wulff, D. U., Meier, D. S., & Mata, R. (2023). Using novel data and ensemble models to improve automated labeling of sustainable development goals. <https://arxiv.org/abs/2301.11353>. [arXiv:2301.11353](https://arxiv.org/abs/2301.11353) [cs]
- Yao, R., Tian, M., Lei, C. U., et al. (2024). Assigning multiple labels of sustainable development goals to open educational resources for sustainability education. *Education and Information Technologies*. <https://doi.org/10.1007/s10639-024-12566-6>. <https://link.springer.com/10.1007/s10639-024-12566-6>
- Yu, Y., Scheidegger, S., Elliott, J., et al. (2024). climatebug: A data-driven framework for analyzing bank reporting through a climate lens. *Expert Systems with Applications*, 239, Article 122162.
- Zhou, K., Liu, Z., Qiao, Y., et al. (2022). Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4), 4396–4415.