

Shared Task

Automatic Classification of the United Nations' Sustainable Development Goals (SDGs) and Their Targets in English Scientific Abstracts

Simon Clematide, Martin Volk, Tobias Fankhauser, Lorenz Hilty and Jürgen Bernard

UZH

Description

In the current era of global challenges, the United Nations' Sustainable Development Goals (SDGs) serve as a universal call to action, addressing critical issues such as poverty, inequality, climate change, environmental degradation, and peace and justice. The alignment of scientific research with these goals is pivotal for measuring and enhancing the impact of academia on these global objectives.

The Swisstext 2024 Shared Task proposed by members of the Departments of Computational Linguistics and Informatics at the University of Zurich targets an innovative and significant undertaking: the automatic classification of scientific abstracts (in English) with respect to SDGs and their specific targets. This task not only aligns with the increasing importance of interdisciplinary research towards sustainable development but also represents a crucial step in integrating advanced Natural Language Processing (NLP) within the realm of sustainability research. By automating the classification process, this Shared Task aims to facilitate the identification and analysis of research towards the SDGs, thereby fostering a targeted and efficient approach in addressing the world's most pressing challenges.

This Shared Task evolved in the context of the UZH Digital Society Initiative project "SDG Research Scout" financed by the Digitalization Initiative of the Zurich Higher Education Institutions (DIZH).

Task 1: Classification at the Level of the 17 SDGs

Motivation

The primary aim of this Shared Task is to evaluate the capability of automated systems in classifying a scientific abstract under the most appropriate SDG. A particular emphasis is placed on the system's proficiency in identifying the best fitting SDG when the specific goals are underrepresented in the dataset, thus ensuring a balanced and comprehensive understanding of the research landscape in relation to the SDGs.

Objective

This Shared Task involves the classification of scientific abstracts based on their relevance to one of the 17 Sustainable Development Goals (SDGs). The objective is to determine the single, most relevant SDG for each abstract. Additionally, if an abstract does not pertain to any of the SDGs, it should be classified under a distinct ‘non-relevant’ category.

Task Details

SDG Single Label Classification: Participants are required to assign one primary SDG to each abstract. This involves identifying the most relevant SDG or classifying the abstract as non-relevant if it does not align with any SDG.

Training Material

A dataset of at least 400 labeled abstracts from the UZH publication repository Zora (zora.uzh.ch) for SDG classification will be provided, facilitating the development of models to predict the most applicable SDG for each abstract. This dataset is unbalanced regarding the 17 SDGs and the topics represented therein.

Evaluation Metrics

Primary Metric: Accuracy. This will assess the correct prediction of the primary SDG for each abstract.

Secondary Metric: Average F1 Score per SDG. This metric is designed to evaluate the balanced performance across all SDGs, providing insight into the model’s ability to accurately classify abstracts under each goal, including those that occur less frequently in the training data.

Additional Resources

Participants are permitted to use any external resources (such as additional services, query collections, or annotated documents such as the large OSDG-CD dataset) to enhance and improve their systems. However, manual correction of system output of the test data is not allowed. It is mandatory for system description papers to clearly specify any additional resources used.

Test Conditions

The participants will get a test set with abstracts. They have to predict the numeric SDG label. The participants will have 2 days to compute and submit the results.

Task 2: Multi-label Classification at the Level of SDG Targets

Task 2 in the Swisstext 2024 Shared Task delves into the more granular aspect of the United Nations' Sustainable Development Goals (SDGs) by focusing on the multi-label classification of SDG targets. Each of the 17 SDGs is underpinned by specific targets that detail the objectives needed to achieve the respective goal. For instance, SDG 3, "Good Health and Well-being," includes targets like reducing maternal mortality and ending epidemics of certain diseases. This task requires participants to predict the most relevant target(s) for a given abstract, based on its classified main SDG from Task 1. If no specific target is relevant for the SDG, a label "unspecific target" should be predicted.

Motivation

This Shared Task is designed to challenge the systems' abilities in fine-grained classification and to encourage the development of innovative solutions in the field of NLP, particularly in scenarios where extensive labeled data is not available. Accurate prediction of SDG targets in scientific abstracts enhances the specificity and usefulness of research classification in the context of sustainable development, fostering a more nuanced understanding of research contributions towards global goals.

Objective

Participants are tasked with predicting the primary target for an abstract associated with a specific SDG. Additionally, where applicable, they should identify other relevant targets within the same SDG. This task aims to deepen the understanding of how specific aspects of research align with detailed objectives of the SDGs.

Task Details

Primary Target Prediction: Given an abstract and its main SDG, the main task is to identify the most relevant target within that goal.

Additional Targets Prediction: If relevant, participants should also predict any number of additional targets that apply to the same SDG for the given abstract.

If the primary label is "unspecific target", no additional targets should be predicted.

Training and Development Material

No specific training data will be provided for target-level classification.

Participants can use the labeled data from Task 1 to select the SDG.

A development set consisting of 50 abstracts with annotated targets will be available for tuning and testing the approaches.

Approach

Given the limited amount of development data, participants are encouraged to explore zero-shot or few-shot learning techniques, particularly prompt engineering, to effectively approach this task. These advanced NLP methodologies are anticipated to be central in participants' solutions due to the task's nature and the data constraints.

Evaluation Metric

The F1 Score will be used to evaluate multi-label target predictions in Task 2. This metric effectively balances precision and recall, assessing model accuracy through true positives, false positives, and false negatives against a manually corrected ground truth.

Test Conditions

The participants will get a test set with abstracts and a provided SDG. They have to predict the main numeric target label and any number of additional target labels for the given SDG of the abstract. The participants will have 2 days to compute and submit the results.

Schedule, System Descriptions and Planned Workshop

System Description Papers

System description papers by participants are required in order to gain insights into which approaches work best for this problem.

SwissText Workshop “NLP for Sustainable Development Goals Monitoring”

A 2-hour workshop, allowing both online and on-site participation, is planned. This workshop aims to bring together task participants and interested stakeholders to delve into the workings of the participating systems. It will focus on identifying the strengths and weaknesses of the proposed approaches, fostering a collaborative environment for open discussion. The workshop will also be a platform to discuss future directions for the research community, targeting enhancements in the automatic classification of SDG-related aspects in scientific work. Attendees will have the opportunity to engage in knowledge sharing, contributing to a collective understanding and shaping the next steps in this application area of research.

Schedule

1. Welcome + Shared-Task Overview: Simon Clematide + Tobias Fankhauser
(15-20 min)
2. Jürgen Bernard (UZH): The DZI Project “SDG Research Scout” and its
visualisation aspects (20 min)
3. - Short Break -
4. Presentation of the Shared Task contributions (15 min each)
 1. Fernando de Meer Pardo, Hanna Hubarava and Vera Bernhard
(UZH): System Description Paper for SwissText 2024 Shared Task
1: Classification at the Level of the 17 SDGs
 2. Manuel Bolz, Andreas Loizidis, Kevin Bründler (UZH): SwissText
- Shared Task (SDG Classification) - Task 1
 3. Adrian M.P. Brasoveanu, Albert Weichselbraun, Lyndon J.B. Nixon
and Arno Scharl (MODUL
University Vienna/webLyzard): An Efficient Workflow Towards
Improving Classifiers in Low-Resource Settings with Synthetic Data
 4. Norman Süsstrunk, Albert Weichselbraun, Andreas Murk, Roger
Waldvogel and André Glatzl (Chur): Scouting out the Border:
Leveraging Explainable AI to Generate Synthetic Training Data
for SDG Classification
5. Final discussion «NLP for Sustainability» (15 min)