

# Towards Robust ESG Analysis Against Greenwashing Risks: Aspect-Action Analysis with Cross-Category Generalization

Keane Ong<sup>♣</sup>, Rui Mao<sup>♣</sup>, Deeksha Varshney<sup>♣</sup>, Erik Cambria<sup>♣</sup>, Gianmarco Mengaldo<sup>♣\*</sup>

<sup>♣</sup>National University of Singapore, Singapore

<sup>♣</sup>Nanyang Technological University, Singapore

{keane.ongweiyang}@u.nus.edu; {deeksha2,mpegim}@nus.edu.sg;

{rui.mao,cambria}@ntu.edu.sg

## Abstract

Sustainability reports are key for evaluating companies' environmental, social and governance (ESG) performance. To analyze these reports, NLP approaches can efficiently extract ESG insights at scale. However, even the most advanced NLP methods lack robustness against ESG content that is greenwashed – i.e. sustainability claims that are misleading, exaggerated, and fabricated. Accordingly, existing NLP approaches often extract insights that reflect misleading or exaggerated sustainability claims rather than objective ESG performance. To tackle this issue, we introduce A3CG - Aspect-Action Analysis with Cross-Category Generalization, as a novel dataset to improve the robustness of ESG analysis amid the prevalence of greenwashing. By explicitly linking sustainability aspects with their associated actions, A3CG facilitates a more fine-grained and transparent evaluation of sustainability claims, ensuring that insights are grounded in verifiable actions rather than vague or misleading rhetoric. Additionally, A3CG emphasizes cross-category generalization. This ensures robust model performance in aspect-action analysis even when companies change their reports to selectively favor certain sustainability areas. Through experiments on A3CG, we analyze state-of-the-art supervised models and LLMs, uncovering their limitations and outlining key directions for future research.

## 1 Introduction

Sustainability reports have become an important mechanism for evaluating a company's environmental, social, and governance (ESG) performance (Nguyen, 2020). To facilitate the analysis of these reports, NLP methods typically offer an automated and efficient way to extract ESG insights at scale (Ong et al., 2024). However, the rise of greenwashing – the practice of providing misleading, exaggerated, and fabricated sustainability

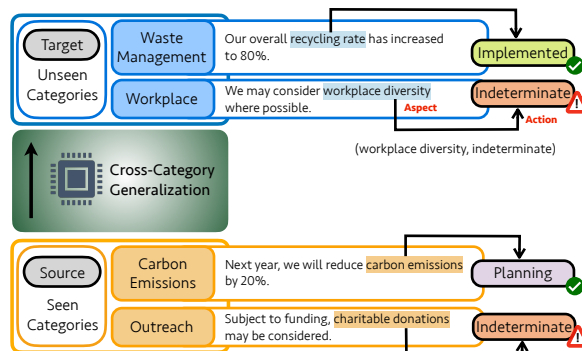


Figure 1: A3CG Task Overview. Given a sustainability statement, the model extracts aspect-action pairs. To evaluate cross-category generalization, it is tested on statements containing aspect categories not seen during training – i.e. trained on "source", tested on "target".

claims – has increasingly obscured the content of these reports (Ong et al., 2024). Greenwashing presents a major challenge for NLP-driven ESG analysis, let alone manual assessments, ultimately hindering meaningful progress toward sustainability goals (Rajesh, 2020).

Yet, existing NLP methods for ESG analysis – i.e. topic analysis (Ong et al., 2025), retrieval-augmented generation (RAG) (Zou et al., 2025), fail to account for greenwashing risks. These approaches aim to provide actionable insights from sustainability reports, but fail to account for the credibility of the claims within them. Therefore, the insights extracted by these NLP methods can often reflect vague and misleading rhetoric instead of genuine sustainability initiatives (Ong et al., 2024).

To tackle this, we propose A3CG<sup>1</sup> - Aspect-Action Analysis with Cross-Category Generalization, a novel dataset for robust ESG analysis amid the prevalence of greenwashing. A3CG provides a foundation for systematically extracting two interconnected components from

<sup>1</sup>We release the dataset and experimental codes at: [https://github.com/keanepotato/a3cg\\_greenwash](https://github.com/keanepotato/a3cg_greenwash)

\*Corresponding author

sustainability texts: *i) Aspect*: The sustainability-related entity, goal, or activity discussed (i.e. “carbon emissions reduction”). *ii) Action*: The type of engagement related to the aspect, “implemented”, “planning”, or “indeterminate”. While “implemented” and “planning” indicate concrete steps taken or planned toward the sustainability aspect, “indeterminate” highlights vague, evasive or non-attributable claims which may be linked with greenwashing (Siano et al., 2017). By enabling the identification of sustainability aspects and their action assessments, A3CG allows models to clarify genuine sustainability initiatives from ambiguous or misleading rhetoric. Through this flagging of clear sustainability initiatives, stakeholders can then verify and track them more easily. Therefore, while A3CG is not designed for *definitive* greenwashing detection, it facilitates *clear* and *transparent* ESG analysis that is less susceptible to greenwashing risks.

To further enhance the robustness of ESG analysis, A3CG extends aspect-action analysis to a cross-category generalization setting. This ensures that aspect-action analysis remains effective even when companies change their report content to selectively favor certain sustainability areas (Darnall et al., 2022). By emphasizing the extraction of aspect-action pairs from sustainability categories beyond those seen in training, A3CG facilitates aspect-action analysis on new, previously unencountered sustainability themes. As a result, aspect-action analysis remains robust in distinguishing genuine sustainability initiatives from rhetoric, even with changes in report content.

We conduct extensive experiments on A3CG, evaluating the performance of state-of-the-art (SOTA) supervised models and large language models (LLMs). Our results highlight: (1) Supervised learning methods - GRACE (Luo et al., 2020), outperform the latest LLMs - Claude 3.5 Sonnet (Huang et al., 2024b), DeepSeek V3 (Liu et al., 2024), on A3CG. (2) Contrastive learning is relatively more effective than adversarial learning for cross-category generalization. (3) The limitations of supervised models and LLMs in A3CG subtasks and linguistic cases. (4) Research directions and strategies for tackling these limitations.

Our main contributions include: (1) We develop and release A3CG, the first dataset tailored for ensuring robust ESG analysis amid the prevalence of greenwashing risks. (2) We conduct extensive experiments on A3CG using SOTA supervised learn-

ing and LLM methods, uncovering valuable insights for NLP research and model development.

## 2 Related Work

**NLP Methods in Sustainability Analysis.** To automate the extraction of insights from sustainability reports, popular NLP approaches include analyzing topic occurrence (Ong et al., 2025), sentiment analysis (Song et al., 2018), and RAG (Zou et al., 2025). Yet, these traditional NLP approaches do not consider how greenwashing might distort extracted insights (Ong et al., 2024). To promote the transparency of sustainability claims, Stambach et al. (2022) studied environmental claim detection, and Ni et al. (2023) assessed report conformance to TCFD guidelines. However, these works do not explicitly address building robust ESG analysis against greenwashing risks.

**Aspect-Based Sentiment Analysis.** Aspect-Based Sentiment Analysis (ABSA) has become a popular NLP problem over the past decade (Du et al., 2023; Ong et al., 2023). Yet, ABSA has focused on reviews and opinion mining (He et al., 2023), despite calls for aspect-level analysis to be extended to other contexts (Chebolu et al., 2023). Our study presents the first adaptation of aspect-level analysis to the aspect-action classification task. While we focus on sustainability, aspect-action analysis could potentially be adapted for other applications - i.e. government policy analysis (Howlett, 2015).

**Cross-Generalization Studies.** Cross-domain generalization has become an increasingly important challenge in NLP (Wang et al., 2022). However, there has been limited investigation on how models perform on *unseen categories within the same domain* - i.e. cross-category generalization. Yet, this deserves attention because intra-domain data shifts are distinct from cross-domain shifts, and can significantly degrade model performance (Subaswamy et al., 2021). Moreover, data annotation for unseen categories is expensive and impractical, as these categories often arise unexpectedly from evolving text content (Bai et al., 2024).

**LLMs Biases in Sustainability Analysis.** Given LLMs’ reasoning capabilities (Wei et al., 2022a), they have increasingly been utilized for tasks within sustainability analysis, such as knowledge base construction (Ong et al., 2025), RAG (Zou et al., 2025), among others. However, their performance on domain-specific tasks can be impacted by pre-training biases (Dai et al., 2024; Yang et al., 2024b),

and the impact of these biases on sustainability analysis tasks remains under-explored.

### 3 Dataset Construction

This section describes A3CG’s construction.

#### 3.1 Data Collection & Quality Control

We collect sustainability statements from 1679 sustainability reports<sup>2</sup> of Singapore Exchange (SGX) companies, from 2017-2022 inclusive. The reports are public and can be found online via company websites. For quality control, we eliminate: (1) Incomplete statements (2) Incoherent statements with misspelled words (3) Non-English statements.

#### 3.2 Annotation Scheme

Pairs of (aspect, action) are annotated for each sustainability statement. In a pair, the *action* characterizes the company’s engagement with the *aspect*. Notably, the aspect’s sustainability category (i.e. "emissions control" as a category of the aspect "carbon emissions") is also annotated. However, similar to other generalization studies (Xu et al., 2023a), the aspect category is not a prediction target. Instead, it is utilized to split the dataset for cross-category testing. We summarize the aspect and action definitions<sup>3</sup> below:

**Aspect:** A sustainability aspect is defined as a significant *entity*, *goal*, *sub-area*, *activity* of a sustainability category, providing a focal point for an action to address or engage with. We focus on aspects that belong to the sustainability categories in Table 5, and are explicitly found within a statement.

**Action:** The type of action taken toward the aspect. *i) Planning:* Indicates that an action has been planned or a commitment has been made by a company to address or engage with the *aspect* to advance its sustainability efforts. *ii) Implemented:* Indicates that an action has already been taken to address or engage with the *aspect* to advance its sustainability efforts. *iii) Indeterminate:* Indicates that it is unclear from the statement if the company intends to address or engage with the *aspect* to advance its sustainability efforts, or how it intends to do so.

<sup>2</sup>Company names and the corresponding years of sustainability reports are also available at: [https://github.com/keanepotato/a3cg\\_greenwash](https://github.com/keanepotato/a3cg_greenwash). Using this information, these public reports can be searched and accessed online.

<sup>3</sup>Detailed annotation instructions, guidelines, definitions, samples, and annotator details can be found in appendix A

For i) and ii), “addressing” or “engaging with” an aspect involves incorporating the aspect within company operations if it is a sustainability *activity* or *entity*, or *advancing* the aspect if it is a sustainability *goal* or *sub-area*. On the other hand, iii) characterizes vague, non-committal and non-attributable language. These action labels differentiate planned commitments from implemented actions, while separating them from ambiguous, non-committal claims.

#### 3.3 Annotation Process

The annotation process involves 5 annotators<sup>3</sup> and 3 verifiers<sup>3</sup>, all actively engaged in doctoral or post-doctoral research in sustainability. (1) *Train & Trial:* Annotators and verifiers undergo rounds of trial annotation, with each round comprising 50 random samples. After each round, annotations are scrutinized for accuracy and conformance to the guidelines, and feedback is provided. The trials occur until each person attains a proficiency of at least 95% correctly labeled samples. (2) *Daily Annotation:* After the trials, annotators label the data daily, and are instructed to flag samples with uncertain annotation. (3) *Resolving Disagreements:* Every 3 days, uncertain annotations are discussed among all annotators to reach a decision. In cases where complete agreement cannot be reached, majority voting is taken. (4) *Validation:* Every 3 days, 20% of annotated data for each annotator (comprising samples that are not flagged with uncertain annotation) are scrutinized by the verifiers for accuracy and conformance to guidelines. When incorrectly annotated statements exceed 5%, annotations for the 3-day period are redone.

#### 3.4 Dataset Statistics

A3CG comprises 2004 statements (samples), comparable to standard aspect-level analysis datasets - Rest15 (Pontiki et al., 2015). Table 1 summarizes the dataset. Moreover, A3CG follows standard aspect-level analysis datasets - Lapt14, Rest14, 15, 16 (Pontiki et al., 2014, 2015, 2016), by having a similar proportion of samples (33.2%) without aspect-action pairs. This adds the challenge of distinguishing sustainability statements that contain aspect-action pairs from those that do not, simulating real-world statements where relevant aspects are not always mentioned (Khan et al., 2016).

Sustainability Category	
Resource Optimisation	368
Emissions Control	284
Worker & Consumer Safety Compliance	328
Workplace	401
Outreach	340
Management	297
Business Compliance	167
Waste Management	307
Data and Cybersecurity Protection	129
Ecological Conservation	102
<b>Total No. Aspects</b>	666
<b>Total Aspects*</b>	2723
Action Type	
Indeterminate	885
Implemented	1459
Planning	379
<b>Total Actions*</b>	2723

Table 1: Aspects and action totals for each sustainability category. \* Totals do not include no-aspect counts.

## 4 Experiments

### 4.1 Experimental Setups

**Full Dataset**<sup>4</sup>. As a preliminary step to verify dataset stability before evaluating cross-category generalization, we split the entire dataset into train, validation, and test sets while keeping a balanced ratio of aspect categories.

**Cross-Category Generalization**<sup>4</sup>. Different from leave-one-domain-out setups in cross-domain generalization (Xu et al., 2023b), cross-category generalization requires a distinct approach. Real-world sustainability statements are complex (Smeuninx et al., 2020), and aspects from different categories tend to co-occur within the same statement, making a leave-one-category-out strategy unrealistic. Therefore, we split the dataset into three equal folds. In each fold, samples are assigned to the training, validation, or test set based on the aspect categories they contain. 3-4 categories are excluded from the training and validation sets, forming the unseen (US) test set. Models are evaluated on the US test set to analyze their capacity for cross-category generalization. A control seen (S) test set is also constructed, comprising categories that overlap with those in the training set, but with entirely different samples. To ensure proper evaluation, no training

or validation sample appears in any test set. The set of excluded categories for each fold is chosen considering their tendency to co-occur in a statement and varies across folds with no overlap, enabling balanced evaluation on different unseen categories.

### 4.2 Models

**Vanilla Models**<sup>4</sup>: Popular models for aspect analysis are utilized, encompassing generative - T5 (Raffel et al., 2023), and sequence tagging - BERT Sequence Tagging, BERT-ST (Mao and Li, 2021).

**Vanilla + Learning Paradigm**<sup>4</sup>: Contrastive learning (CL) and adversarial learning (AL) are applied to vanilla models - T5+CL, T5+AL, BERT-ST+CL, BERT-ST+AL. CL uses the supervised contrastive loss for pre-training (Li et al., 2021):  $\mathcal{L}_{\mathcal{I}}^{\text{sup}} = \sum_{i \in \mathcal{I}} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_p / \tau)}{\sum_{a \in N(i)} \exp(\mathbf{z}_i \cdot \mathbf{z}_a / \tau)}$ . In batch  $\mathcal{I}$ , index  $i$  represents the anchor sample.  $P(i) = \{p \in \mathcal{I} : p \neq i\}$  is the set of indices of all positives distinct from  $i$  (samples with the same category label as the anchor; a sample’s category label is determined by the categories of its contained aspect(s)).  $N(i) = \{n \in \mathcal{I} : n \notin P(i)\}$  is the set of indices of all negatives (samples not sharing any category label with the anchor). Following Ganin and Lempitsky (2015), AL focuses on learning category-invariant features. A category discriminator is added to the encoder to predict the categories of a sample, with a gradient reversal layer to reverse the discriminator loss gradients.

**SOTA-ABSA**<sup>4</sup>: We adapt SOTA-ABSA models for A3CG, given that both ABSA and A3CG tasks overlap in their focus on aspect-level analysis. Therefore, we hypothesize that ABSA methods may be applicable to A3CG, and test this through experiments. We select the latest sequence-tagging methods—GRACE (Luo et al., 2020), and generative methods—InstructABSA (Scaria et al., 2023), IT-RER-ABSA (Zheng et al., 2024), CONTRASTE (Mukherjee et al., 2023).

**Large Language Models**<sup>4</sup>: The latest LLMs (zero-shot and few-shot) are evaluated. These are selected from different LLM families, and have achieved state of the art performance on general reasoning and natural language understanding benchmarks. The LLMs evaluated include—GPT-4o (OpenAI, 2024), Claude 3.5 Sonnet (Huang et al., 2024b), Llama 3 (70B) (AI@Meta, 2024), DeepSeek V3 (Liu et al., 2024).

<sup>4</sup>Appendix C details all experimental setups and model implementations, including hyperparameter study for AL.



Method	Full Dataset	Fold 1		Fold 2		Fold 3		S Avg	US Avg	$\Delta$
		S	US	S	US	S	US			
T5	<u>70.48</u>	<u>57.85</u>	<u>43.03</u>	<u>68.90</u>	<u>45.74</u>	<u>67.94</u>	<u>34.59</u>	<u>64.90</u>	<u>41.12</u>	-23.78
BERT-ST	43.19	39.56	25.25	39.00	22.92	43.97	30.01	40.84	26.06	-14.78
T5 + CL	<u>71.12</u>	<b>62.96</b>	<u>46.97</u>	<u>69.76</u>	<u>46.67</u>	<u>67.99</u>	<u>38.33</u>	<u>66.90</u>	<u>43.99</u>	-22.91
T5 + AL	69.27	61.24	39.62	66.91	<u>47.02</u>	65.17	<u>41.82</u>	64.44	42.82	-21.62
BERT-ST + CL	68.53	60.00	37.06	69.22	41.78	58.04	34.94	62.42	37.93	-24.49
BERT-ST + AL	24.30	37.75	27.05	27.13	25.57	34.70	26.11	33.19	26.24	-6.95
InstructABSA	69.47	60.23	37.53	64.73	49.76	64.14	47.38	63.03	44.89	-18.14
IT-RER-ABSA	69.20	57.70	41.81	66.02	48.87	68.83	39.10	64.18	43.26	-20.96
GRACE	67.09	60.87	<b>50.08</b>	63.10	44.33	61.92	<b>48.12</b>	61.96	<b>47.51</b>	-14.45
CONTRASTE	<b>71.33</b>	<u>61.26</u>	<u>48.14</u>	<b>69.81</b>	<b>50.30</b>	<b>71.34</b>	<u>40.58</u>	<b>67.47</b>	<u>46.34</u>	-21.13
GPT-4o	29.79	31.61	42.98	40.00	32.51	35.58	32.35	35.73	35.95	+0.22
GPT-4o + FS*	35.69	39.08	46.68	39.12	41.10	40.05	33.46	39.42	40.41	+0.99
Claude 3.5 Sonnet	37.70	36.71	39.44	41.11	36.88	41.59	38.22	39.80	38.18	-1.62
Claude 3.5 Sonnet + FS*	<u>42.11</u>	<u>40.62</u>	<u>46.27</u>	<u>43.18</u>	<u>40.35</u>	<u>45.04</u>	<u>39.48</u>	<u>42.95</u>	<u>42.03</u>	-0.92
Llama 3 (70B)	<u>20.15</u>	17.97	25.24	23.66	18.33	22.31	18.43	21.31	20.67	-0.64
Llama 3 (70B)+ FS*	29.82	25.03	38.11	33.49	33.30	33.15	26.65	30.56	32.69	+2.13
DeepSeek V3	40.16	38.06	<u>48.18</u>	<u>45.27</u>	40.23	42.46	34.84	41.96	41.08	-0.88
DeepSeek V3 + FS*	35.63	38.85	30.16	44.18	<u>44.51</u>	<u>46.62</u>	34.43	<u>43.22</u>	36.37	-6.85

Table 2: Aspect-Action Analysis (AAA) F1 across full dataset and folds with seen (S) and unseen (US) categories. \*FS denotes few-shot, with results averaged over 5 random samplings for the few shot examples.  $\Delta$  represents US Avg - S Avg. Best results per model type (following section 4.2) are underlined; overall best results are bolded.

## 5 Results & Discussion

From Table 2, model evaluation on the full dataset indicates the dataset’s stability, with supervised models (T5, T5+CL, CONTRASTE) achieving over 70% F1. In the following, we focus on A3CG’s core objective: aspect-action analysis (AAA) - the extraction of aspect-action pairs in a cross-category generalization setup. Therefore, we primarily discuss performance on unseen (US) categories across all folds. Performance on sub-tasks of AAA is also analyzed: aspect term extraction (ATE) for identifying aspects, and action classification (AC) for classifying action labels after correctly identifying aspects (AC excludes no-aspect cases).

For the models evaluated in our paper, we distinguish between two types of models: supervised models and large language models (LLMs). Supervised models include, from Section 4.2—vanilla models, vanilla + learning paradigm and SOTA-ABSA. These models are based on pre-trained language models (PLMs) (i.e. T5, BERT), which are typically smaller in scale and fine-tuned for specific tasks (in our case, the A3CG task). In contrast, when we refer to LLMs, we refer to larger, general-purpose models (i.e. GPT-4o) that are trained on broad language tasks, and are implemented here via prompt-based methods without any fine-tuning. In the following, we discuss the results of our experiments, by first analyzing the performance of supervised methods (Sections 5.1, 5.2), followed by

LLMs (Sections 5.3, 5.4). At the end, we compare both supervised methods and LLMs (Section 5.5) to contrast the strengths and weaknesses of both methods for A3CG.

### 5.1 A3CG Challenges for Supervised Models




Figure 2: Supervised models’ drop in ATE precision, recall, F1, from seen (S) to unseen (US) categories, averaged across all folds (US Avg score - S Avg score).

**Supervised models struggle to detect aspects of unseen categories.** From Figure 2, ATE recall drops significantly from seen to unseen categories, averaged across folds. Across all supervised models, the recall decline is notably greater than the precision decline, suggesting that the drop in F1 scores is considerably driven by lower recall. Qualitative




Figure 3: Average F1 score across all folds of unseen categories (US Avg) for supervised and LLM methods. AAA is the main task of A3CG, ATE and AC are subtasks of AAA. For detailed breakdown, refer to Table 14.

analysis indicates that supervised models cannot generalize to sustainability aspects that, while semantically dissimilar to aspects in seen categories, are still highly intuitive to sustainability (i.e. “carbon emissions”). From Table 12, these obvious sustainability aspects are typically undetected in ATE. This suggests that models potentially rely more on category-specific patterns instead of a broader conceptual understanding of sustainability-related aspects. Therefore, *to improve generalization, supervised models can be enhanced by incorporating contextually-aware representations (Liang et al., 2022) that capture broader sustainability concepts.*

**Qualitative analysis reveals that supervised models primarily struggle with specific cases of complex syntax.** These prominent failure cases include: i) *Elipsis*: Misclassifying actions when key words that explicate the action taken on the aspect are omitted, with contextual inference required for action classification instead. ii) *Ambiguous Syntax*: When it is unclear how a phrase modifies an aspect, creating ambiguity in the action taken, the model misclassifies the action as a concrete implementation or plan instead of “indeterminate”. Table 12 shows these examples. Hence, *to improve performance on complex syntax cases, supervised models can incorporate syntactic-aware architectures (Huang et al., 2024a) to enhance syntactic contextual inference and resolve ambiguous syntax.*

## 5.2 Comparison of Learning Paradigms for Supervised Models

Popular methods for cross-domain generalization, adversarial learning (AL) and contrastive learning

(CL), are evaluated on cross-category generalization within A3CG. From Figure 4(a), compared to vanilla T5 and BERT-ST baselines, although AL still enhances generalization performance (+1.70 for T5+AL; +0.18 for BERT+AL), CL yields relatively higher F1 improvements on unseen categories (+2.87 for T5+CL; +11.87 for BERT+CL), averaged across all folds.

The generalization effectiveness of CL could be attributed to the models’ ability to learn robust semantic features (category distinctions) that are transferable across categories (Khosla et al., 2020). In a case study, the feature embeddings of samples from different unseen categories exhibit more distinct clustering for Figure 4(c) T5+CL, compared to 4(b) T5 Vanilla, which shows greater overlap. This suggests that CL allows the model to capture semantic distinctions between unseen categories, reducing category overlap and ambiguity. Therefore, instead of spurious correlations, the model relies more on robust semantic features unique to each category, improving unseen category generalization (Izmailov et al., 2022).

In contrast, AL focuses on learning category-invariant features induced by feature collapse to promote generalization (Tang and Jia, 2020). From Figure 6(d) T5+AL, features from samples of unseen categories appear less discriminable and more entangled, particularly along the y-axis, compared to 6(b) T5 Vanilla. However, by collapsing features for invariance, AL may inadvertently suppress domain-specific feature attributes that are relevant for identifying aspects of unseen categories. Hence, though AL increases generalization performance,

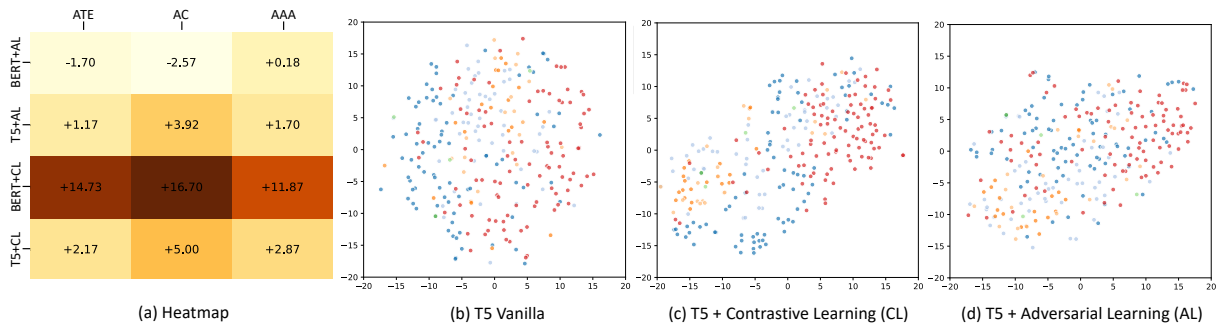


Figure 4: (a) Heatmap of F1 improvements on unseen categories, averaged across all folds, for vanilla+learning paradigm models, compared to original vanilla models. (b), (c), (d) t-SNE visualization of feature embeddings for fold 3 unseen category samples, with samples color-coded by category. Figure 6 shows t-SNE with a full legend.

it yields lower improvements than CL.

This presents a key consideration for supervised training techniques in A3CG: *although AL is still beneficial, CL can be a more effective strategy for unseen category generalization.*

### 5.3 Comparison of LLM performance

**LLMs with higher reasoning capabilities perform better on A3CG.** From Table 2, Claude 3.5 Sonnet+FS and DeepSeek V3 yield the highest average F1 on unseen categories among LLMs (US Avg 42.03 & 41.08 respectively). This aligns with their superior performance over the other LLMs - GPT-4o and Llama 3 (70B) on general reasoning benchmarks (Liu et al., 2024; Dubey et al., 2024). This suggests that A3CG goes beyond naive text classification to require structured reasoning. Therefore, *selecting LLMs with higher reasoning capabilities can optimize performance on A3CG.*

**Few-shot improves performance on nearly all LLMs except DeepSeek V3.** From Figure 3, the addition of few-shot examples to LLMs generally improves their F1 performance on AAA for unseen categories. This suggests that LLMs tend to generalize better to unseen categories through in-context learning (Dong et al., 2024). Yet, including few-shot examples for DeepSeek V3 unexpectedly degrades model performance. Preliminary studies on a related model - DeepSeek R1 (Guo et al., 2025) highlight its over-reliance and failure to generalize from examples (Parmar and Govindarajulu, 2025). Yet, these studies may only have limited applicability to DeepSeek V3, warranting further investigation. Therefore, *to optimize performance on A3CG, few-shot examples can be provided for LLMs such as GPT-4o, Claude 3.5 Sonnet, Llama 3 (70B). In contrast, DeepSeek V3 unexpectedly performs better in zero-shot, underscoring the need*

*for further studies into its few-shot capabilities.*

### 5.4 A3CG Challenges for LLMs

**LLMs generally exhibit higher recall toward aspects of environmental categories, potentially due to pre-training biases.** Figure 5 shows that ATE recall on all environmental categories (*resource optimization, emissions control, ecological conservation*) tend to exceed the median ATE recall of non-environmental categories, frequently surpassing their upper quartile. This suggests that LLMs capture environmental aspects more effectively than non-environmental ones, potentially due to popularity biases in their pre-training (Dai et al., 2024). Since LLMs are pre-trained on large-scale user-generated data, their learned distributions are often influenced by the popularity of specific topics - in this case, the prominent association of sustainability with environmental protection over non-environmental areas (i.e. social and governance) (Ruggerio, 2021). Consequently, LLMs may end up disproportionately favoring environmental aspects for ATE while under-representing non-environmental sustainability aspects. However, while our findings highlight this possibility, further research is needed to assess how much these biases contribute to the recall skew. Therefore, *to mitigate this imbalance for improved performance, methods such as bias-aware re-ranking (Carraro and Bridge, 2024) can be employed to balance the selection of non-environmental aspects, although deeper investigation into the causes of the recall skew could enhance mitigation strategies.*

**Qualitative analysis reveals that LLMs primarily struggle with specific cases of complex pragmatics.** Pragmatics deals with the implicit elements of language that convey intent and contextual nuance beyond literal meaning (Mao et al., 2025).

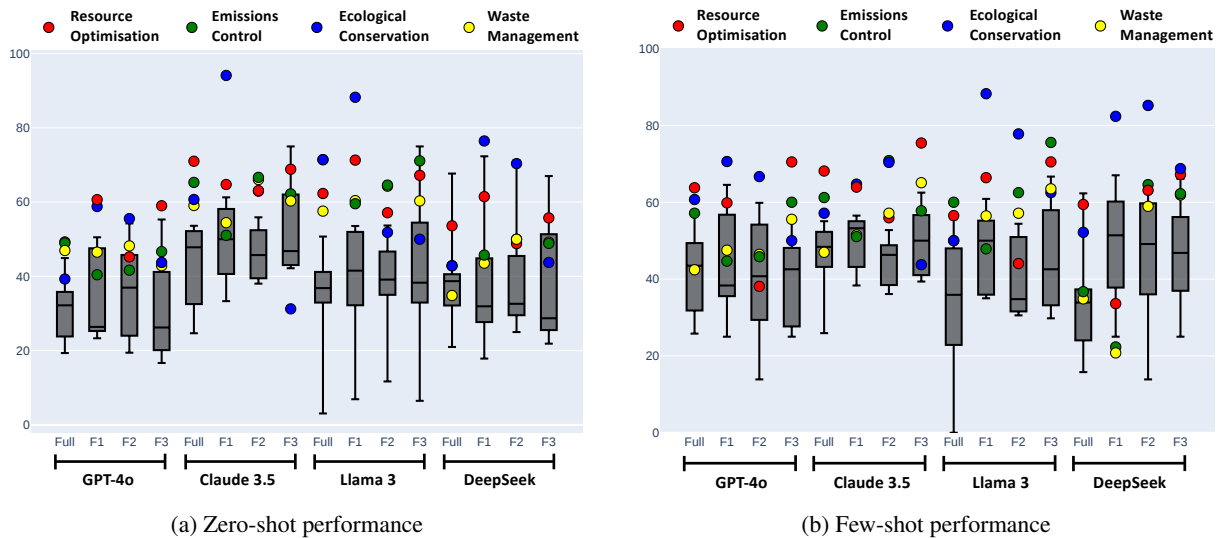


Figure 5: ATE recall scores for categories in the Full Dataset (Full) test set and Fold (F1, F2, F3) seen & unseen test sets. Environmental categories (circles) generally have greater recall than non-environmental categories (box plots).

In both zero-shot and few-shot settings, LLMs’ most prominent failure cases involve pragmatic challenges that affect the intent, commitment, and certainty of sustainability claims: i) *Misinterpreting modality*: Misclassifying actions as concrete commitments or definite plans, despite the statement containing tentative and hedging language that indicate uncertainty. ii) *Negation handling*: Misclassifying negative actions as implementations or plans to improve sustainability. iii) *Unattributed statements*: Misclassifying actions as concrete company-specific implementations or plans, despite the statement lacking explicit attribution to the company. iv) *Future Dependency*: Misclassifying planned actions as already implemented, despite the statement emphasizing a future dependency required for the action to occur. Table 13 details these examples. Therefore, *to adapt LLMs for capturing pragmatic cues, prompting strategies such as chain-of-thought (Wei et al., 2022b) and self-reflection (Ji et al., 2023), can be employed to explicitly reason through modality, negation, attribution, future dependency before classification, thereby improving robustness on these pragmatic cases.*

## 5.5 Comparison of Model Types

**Supervised models can perform better than LLMs on unseen categories**, with GRACE achieving the highest average F1 score on unseen categories (US Avg 47.51) compared to LLMs, from Table 2. As shown in Figure 3, supervised models generally outperform LLMs in AAA on unseen cat-

egories, except for T5, BERT-ST, BERT-ST+CL, BERT-ST+AL. This suggests that while LLMs perform well on a broad range of tasks, supervised fine-tuning is generally better at capturing task-specific patterns within A3CG. Thus, *supervised fine-tuning remains essential for optimal performance on A3CG.*

**While supervised models have higher absolute performance than LLMs, they show significant inefficiency at transferring performance from seen to unseen categories.** For supervised models, although they demonstrate the best absolute performance on unseen categories, the difference between average F1 scores on seen and unseen categories,  $\Delta$ , is significant - greater than 14% in magnitude except BERT-ST+AL ( $\Delta=-6.95\%$ ), from Table 2. This represents a substantial drop in performance when transferring learned patterns from seen to unseen categories, underscoring the weak generalization efficiency of learned features. The features learned by LLMs are shaped by broad pre-training, making them generally applicable across diverse data distributions. As a result, the performance drop on unseen categories is relatively modest, as these categories do not represent truly out-of-distribution data. In contrast, supervised models rely on features tuned to the training distribution, and their performance shortfall on unseen categories reflects a stronger sensitivity to distribution shifts. Consequently, *supervised models can improve the generalization efficiency of their learned features, to enhance cross-category generalization.*



**Specific LLMs can still perform competitively on unseen categories.** Despite being generally outperformed by supervised models, specific LLMs achieve competitive average F1 scores on unseen categories, narrowing the performance gap. From Figure 3, Claude 3.5 Sonnet+FS (42.03) outperforms several supervised models - T5 (41.12), BERT-ST (26.06), BERT-ST+CL (37.93), BERT-ST+AL (26.24), while lagging behind IT-RER-ABSA (43.26), T5+AL (42.82) by less than 1.5%. Therefore, *while supervised learning provides optimal performance, its superiority is not absolute - specific LLMs can still present viable alternatives to full-scale supervised fine-tuning, particularly in scenarios where computational resources are limited but API-based access to LLMs are available.*

**Action classification is more challenging than aspect term extraction for all methods.** From Figure 3, AC F1 scores are markedly lower than ATE on unseen categories across all methods. This is potentially because while both subtasks involve semantic, syntactic and pragmatic reasoning, AC places greater emphasis on interpreting complex syntax and pragmatics, as discussed in the failure cases in supervised models (Section 5.1) and LLMs (Section 5.4). Hence, AC is particularly challenging for supervised and LLM approaches. Therefore, *given the considerable difficulty and lower performance, targeted model improvements for AC could drive significant gains in the overall performance of AAA.*

**Comparative error analysis reveals that LLMs excel in syntax but struggle with pragmatics, while supervised models display the reverse tendency.** Unlike supervised models, LLMs are largely unaffected by syntactic challenges (*ellipsis, ambiguous syntax*) in A3CG, possibly due to their broad syntactic exposure from large-scale pre-training (Yang et al., 2024a). However, relative to supervised models, LLMs struggle with specific pragmatic cases (*modality, negation, unattributed statements, future dependency*) that shape how sustainability claims are framed. This suggests that relying solely on pre-training constrains LLMs' pragmatic reasoning capabilities, limiting their ability to interpret the nuanced, context-dependent discourse in A3CG. Without fine-tuning, LLMs tend to rely on dominant linguistic patterns from pre-training, rather than aligning with the pragmatic cues of a given context (Ruis et al., 2024). This underscores a key point: *while prompting techniques can elicit reasoning in pragmatic cases, they re-*

*main susceptible to patterns in the pre-training distribution. Supervised fine-tuning, in contrast, adjusts the training distribution to more effectively internalize pragmatic cues, allowing more reliable contextual interpretation. This is significant in A3CG, let alone the greenwashing context, where pragmatic cases like hedging, indirect intent, strategic ambiguity, are prevalent (Siano et al., 2017).*

## 6 Conclusion

In this study, we discussed the greenwashing phenomenon, and highlighted significant limitations of current NLP methods for robust ESG analysis. To address this, we proposed the A3CG dataset and evaluated supervised models, and LLMs under zero and few-shot settings. While supervised models achieved the highest absolute performance for unseen category generalization, LLMs exhibited greater generalization efficiency from seen to unseen categories. Within supervised approaches, contrastive learning proved more effective than adversarial learning for unseen category generalization. Yet both supervised and LLM approaches face limitations. On the one hand, supervised models lack contextual awareness for detecting aspects of unseen categories and struggle to classify actions in complex syntactic cases. On the other hand, LLMs overlook non-environmental aspects possibly due to pretraining biases, and struggle with classifying actions where there is pragmatic nuance. These results underscore the need for further research into models that can generalize robustly across ESG categories for the aspect-action analysis task. Therefore, we invite researchers to experiment with A3CG to unlock the potential of NLP for robust ESG analysis.

## Limitations

The A3CG dataset focuses on sustainability statements that are solely in English. However, we acknowledge that sustainability is a global responsibility that can encompass many different regions. Therefore, for greater inclusivity, future work will extend A3CG to non-English languages. Additionally, due to computational resources, this paper does not consider supervised fine-tuning of smaller LLMs—i.e. Llama 3 8B (AI@Meta, 2024). Yet, LoRA-based fine-tuning of these smaller scale LLMs has proven effective in other domain-specific tasks (Varshney et al., 2025), and could represent a future research direction.

## Ethical Considerations

Procedures for data collection were approved by an internal ethics review board within our research group. Additionally, we adhere to ethical principles by ensuring that all data collection and processing are performed with respect for privacy and confidentiality. We use publicly available sustainability disclosures from respective company websites. Considering that these disclosures may contain company and personal names, we make efforts to anonymize sensitive and personal information in the A3CG dataset, focusing solely on the sustainability content relevant to our research. Additionally, the models utilized are publicly available, found from published research papers. Our usage of all data, packages and models adheres to the copyright guidelines provided by the respective copyright holders. Human annotators follow strict guidelines to maintain objectivity and reduce bias. We also ensure transparency in our methodology and provide clear attribution for our sources, aiming to support ethical practices in data usage and dissemination. The dataset and experimental code used in this study will be made publicly available upon acceptance, for the sole purpose of facilitating research, accompanied by appropriate copyright provisions.

## Acknowledgements

This research/project is supported by the NUS Sustainable and Green Finance Institute (SGFIN), NUS Asian Institute of Digital Finance (AIDF), Ministry of Education, Singapore under its MOE Academic Research Fund Tier 2 (STEM RIE2025 Award MOE-T2EP20123-0005: “Neurosymbolic AI for Commonsense-based Question Answering in Multiple Domains”), MOE Tier 2 Award (MOE-T2EP50221-0006: “Prediction-to-Mitigation with Digital Twins of the Earth’s Weather”), MOE Tier 1 Award (MOE-T2EP50221-0028: “Discipline-Informed Neural Networks for Interpretable Time-Series Discovery”), and by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026: “Generative AI”), administered by A\*STAR, as well as supported by Alibaba Group and NTU Singapore.

## References

AI@Meta. 2024. Llama 3 model card. <https://github.com/meta-llama/llama3/>

[blob/main/MODEL\\_CARD.md](#). Accessed: [date of access].

Yinhao Bai, Yalan Xie, Xiaoyi Liu, Yuhua Zhao, Zhixin Han, Mengting Hu, Hang Gao, and Renhong Cheng. 2024. Bvsp: Broad-view soft prompting for few-shot aspect sentiment quad prediction. *arXiv preprint arXiv:2406.07365*.

Diego Carraro and Derek Bridge. 2024. Enhancing recommendation diversity by re-ranking with large language models. *ACM Transactions on Recommender Systems*.

Siva Uday Sampreeth Chebolu, Franck Dernoncourt, Nedim Lipka, and Thamar Solorio. 2023. A review of datasets for aspect-based sentiment analysis. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 611–628.

Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. Bias and unfairness in information retrieval systems: New challenges in the llm era. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6437–6447.

Nicole Darnall, Hyunjung Ji, Kazuyuki Iwata, and Toshi H. Arimura. 2022. [Do esg reporting guidelines and verifications enhance firms’ information disclosure?](#) *Corporate Social Responsibility and Environmental Management*, 29:1214–1230.

Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024. [A survey on in-context learning](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, Miami, Florida, USA. Association for Computational Linguistics.

Kelvin Du, Frank Xing, and Erik Cambria. 2023. Incorporating multiple knowledge sources for targeted aspect-based financial sentiment analysis. *ACM Transactions on Management Information Systems*, 14(3):23.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Li Li Eng, Mahelet Fikru, and Thanyaluk Vichitsarawong. 2022. [Comparing the informativeness of sustainability disclosures versus esg disclosure ratings](#). *Sustainability Accounting, Management and Policy Journal*, 13:494–518.

Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR.

- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of machine learning research*, 17(59):1–35.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Kai He, Rui Mao, Tieliang Gong, Chen Li, and Erik Cambria. 2023. Meta-based self-training and re-weighting for aspect-based sentiment analysis. *IEEE Transactions on Affective Computing*, 14(3):1731–1742.
- Michael Howlett. 2015. Policy analytical capacity: The supply and demand for policy analysis in government. *Policy and Society*, 34(3-4):173–182.
- Xiaosai Huang, Jing Li, Jia Wu, Jun Chang, Donghua Liu, and Kai Zhu. 2024a. Flexibly utilizing syntactic knowledge in aspect-based sentiment analysis. *Information Processing & Management*, 61(3):103630.
- Zhen Huang, Zengzhi Wang, Shijie Xia, and Pengfei Liu. 2024b. Olympicarena medal ranks: Who is the most intelligent ai so far? *arXiv preprint arXiv:2406.16772*.
- Pavel Izmailov, Polina Kirichenko, Nate Gruver, and Andrew G Wilson. 2022. On feature learning in the presence of spurious correlations. *Advances in Neural Information Processing Systems*, 35:38516–38532.
- Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. 2023. Towards mitigating hallucination in large language models via self-reflection. *arXiv preprint arXiv:2310.06271*.
- Mozaffar Khan, George Serafeim, and Aaron Yoon. 2016. Corporate sustainability: First evidence on materiality. *The accounting review*, 91(6):1697–1724.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673.
- Zhengyan Li, Yicheng Zou, Chong Zhang, Qi Zhang, and Zhongyu Wei. 2021. Learning implicit sentiment in aspect-based sentiment analysis with supervised contrastive pre-training. *arXiv preprint arXiv:2111.02194*.
- Bin Liang, Hang Su, Lin Gui, Erik Cambria, and Ruifeng Xu. 2022. Aspect-based sentiment analysis via affective knowledge enhanced graph convolutional networks. *Knowledge-Based Systems*, 235:107643.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Huaishao Luo, Lei Ji, Tianrui Li, Nan Duan, and Daxin Jiang. 2020. GRACE: Gradient harmonized and cascaded labeling for aspect-based sentiment analysis. In *Findings of EMNLP*.
- Rui Mao, Mengshi Ge, Sooji Han, Wei Li, Kai He, Luyao Zhu, and Erik Cambria. 2025. A survey on pragmatic processing techniques. *Information Fusion*, 114:102712.
- Rui Mao and Xiao Li. 2021. Bridging towers of multi-task learning with a gating mechanism for aspect-based sentiment analysis and sequential metaphor identification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 13534–13542.
- Rajdeep Mukherjee, Nithish Kannan, Saurabh Kumar Pandey, and Pawan Goyal. 2023. Contrast: Supervised contrastive pre-training with aspect-based prompts for aspect sentiment triplet extraction. *arXiv preprint arXiv:2310.15577*.
- Doan Thi Thuc Nguyen. 2020. An empirical study on the impact of sustainability reporting on firm value. *Journal of Competitiveness*.
- Jingwei Ni, Julia Bingler, Chiara Colesanti-Senni, Mathias Kraus, Glen Gostlow, Tobias Schimanski, Dominik Stammach, Saeid Ashraf Vaghefi, Qian Wang, Nicolas Webersinke, et al. 2023. Chatreport: Democratizing sustainability disclosure analysis through llm-based tools. *arXiv preprint arXiv:2307.15770*.
- Keane Ong, Rui Mao, Ranjan Satapathy, Ricardo Shirota Filho, Erik Cambria, Johan Sulaeman, and Gianmarco Mengaldo. 2024. Explainable natural language processing for corporate sustainability analysis. *Information Fusion*, page 102726.
- Keane Ong, Rui Mao, Frank Xing, Ranjan Satapathy, Johan Sulaeman, Erik Cambria, and Gianmarco Mengaldo. 2025. ESGSenticNet: A neurosymbolic knowledge base for corporate sustainability analysis. *arXiv preprint arXiv:2501.15720*.
- Keane Ong, Wihan van der Heever, Ranjan Satapathy, Erik Cambria, and Gianmarco Mengaldo. 2023. FinXABSA: Explainable finance through aspect-based sentiment analysis. In *2023 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 773–782. IEEE.
- OpenAI. 2024. Hello GPT-4O. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2024-11-23.
- Manojkumar Parmar and Yuvaraj Govindarajulu. 2025. Challenges in ensuring ai safety in deepseek-r1 models: The shortcomings of reinforcement learning strategies. *arXiv preprint arXiv:2501.17030*.



- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Nuria Bel, Salud María Jiménez-Zafra, and Gülşen Eryiğit. 2016. [SemEval-2016 task 5: Aspect based sentiment analysis](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 19–30, San Diego, California. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. [SemEval-2015 task 12: Aspect based sentiment analysis](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 486–495, Denver, Colorado. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Haris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [SemEval-2014 task 4: Aspect based sentiment analysis](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, Dublin, Ireland. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Preprint*, arXiv:1910.10683.
- R Rajesh. 2020. Exploring the sustainability performances of firms using environmental, social, and governance scores. *Journal of Cleaner Production*, 247:119600.
- Carlos Alberto Ruggerio. 2021. Sustainability and sustainable development: A review of principles and definitions. *Science of the Total Environment*, 786:147481.
- Laura Ruis, Akbir Khan, Stella Biderman, Sara Hooker, Tim Rocktäschel, and Edward Grefenstette. 2024. The goldilocks of pragmatic understanding: Fine-tuning strategy matters for implicature resolution by llms. *Advances in Neural Information Processing Systems*, 36.
- Kevin Scaria, Himanshu Gupta, Siddharth Goyal, Saurabh Arjun Sawant, Swaroop Mishra, and Chitta Baral. 2023. [Instructabsa: Instruction learning for aspect based sentiment analysis](#). *arXiv preprint arXiv:2302.08624*.
- Alfonso Siano, Agostino Vollero, Francesca Conte, and Sara Amabile. 2017. “more than words”: Expanding the taxonomy of greenwashing after the volkswagen scandal. *Journal of business research*, 71:27–37.
- Nils Smeuninx, Bernard De Clerck, and Walter Aerts. 2020. [Measuring the readability of sustainability reports: A corpus-based analysis through standard formulae and nlp](#). *International Journal of Business Communication*, 57(1):52–85.
- Yuan Song, Hongwei Wang, and Maoran Zhu. 2018. Sustainable strategy for corporate governance based on the sentiment analysis of financial reports with csr. *Financial Innovation*, 4:1–14.
- Dominik Stammbach, Nicolas Webersinke, Julia Anna Bingler, Mathias Kraus, and Markus Leippold. 2022. Environmental claim detection. *arXiv preprint arXiv:2209.00507*.
- Adarsh Subbaswamy, Roy Adams, and Suchi Saria. 2021. Evaluating model robustness and stability to dataset shift. In *International conference on artificial intelligence and statistics*, pages 2611–2619. PMLR.
- Hui Tang and Kui Jia. 2020. Discriminative adversarial domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 5940–5947.
- Deeksha Varshney, Keane Ong, Rui Mao, Erik Cambria, and Gianmarco Mengaldo. 2025. Climaempact: Domain-aligned small language models and datasets for extreme weather analytics. *arXiv preprint arXiv:2504.19066*.
- Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and S Yu Philip. 2022. Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8):8052–8072.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022a. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Ting Xu, Huiyun Yang, Zhen Wu, Jiase Chen, Fei Zhao, and Xinyu Dai. 2023a. Measuring your llm models in the wild: A diversified multi-domain dataset for aspect sentiment triplet extraction. *arXiv preprint arXiv:2305.17448*.
- Ting Xu, Huiyun Yang, Zhen Wu, Jiase Chen, Fei Zhao, and Xinyu Dai. 2023b. Measuring your llm models in the wild: A diversified multi-domain dataset for aspect sentiment triplet extraction. *arXiv preprint arXiv:2305.17448*.
- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Shaochen Zhong, Bing Yin, and Xia Hu. 2024a. Harnessing the power of llms in practice: A survey on chatgpt and beyond. *ACM Transactions on Knowledge Discovery from Data*, 18(6):1–32.



Nakyeong Yang, Taegwan Kang, Stanley Jungkyu Choi, Honglak Lee, and Kyomin Jung. 2024b. Mitigating biases for instruction-following language models via bias neurons elimination. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9061–9073.

Guangmin Zheng, Jin Wang, Liang-Chih Yu, and Xuejie Zhang. 2024. Instruction tuning with retrieval-based examples ranking for aspect-based sentiment analysis. *arXiv preprint arXiv:2405.18035*.

Yi Zou, Mengying Shi, Zhongjie Chen, Zhu Deng, Zongxiong Lei, Zihan Zeng, Shiming Yang, Hongxiang Tong, Lei Xiao, and Wenwen Zhou. 2025. [ES-GRReveal: An LLM-based approach for extracting structured data from esg reports](#). *Journal of Cleaner Production*, 489:144572.

## A Dataset

### A.1 Dataset Samples

Additional samples from the A3CG dataset are presented in Table 4, covering sustainability categories not shown in Figure 1. This provides an appreciation of samples from all sustainability categories, including samples that comprise aspects from multiple categories.

### A.2 Annotators

There are a total of 5 human annotators and 3 human verifiers, all based in Singapore and recruited from reputable research institutions. These annotators and verifiers are affiliated with the Asian Institute of Digital Finance and the National University of Singapore College of Design and Engineering. The annotators and verifiers are actively engaged in doctoral or post-doctoral research in the corporate sustainability field, with expertise in sustainability analysis. They contribute to the annotation process as part of their formal academic and research activities, and are compensated above the local minimum wage (SGD\$15/hr) for their involvement. They strictly follow the annotation scheme and guidelines while annotating the dataset, and have consented to its use for research purposes.

### A.3 Annotation Instructions

This section details the full text of instructions given to all annotators. Each annotator is instructed to read the background, before the general instructions, aspect and action annotation guidelines.

**Background:** Sustainability reports are important for evaluating a company’s environmental, social

and governance (ESG) performance. These reports are regularly released by companies, and contain sustainability statements that describe the company’s sustainability efforts. For this annotation task, you will be analyzing sustainability statements to annotate them for research purposes.

### General Instructions:

1. Please annotate (aspect, action) pairs within sustainability statements.
2. For each pair, the aspect and actions are connected. This means that the aspect provides a focal point for the action to address or engage with.
3. Annotate the aspect and actions according to the respective aspect and action guidelines that we will provide.
4. Each aspect can be classified under a specific sustainability category that it belongs to. Please output the associated category for each aspect.
5. If there are no aspects within the sustainability statement, annotate the statement with a (“no aspect”, “no action”) label.
6. For sustainability statements where you are uncertain about the annotation, please flag them up for discussion with the rest of the reviewing team.

### Guidelines for Aspect Annotation:

1. An aspect describes a part or attribute of a specific sustainability category, that is explicitly found within the sustainability statement.
2. In our case, an aspect can either be an *entity*, *goal*, *sub-area*, *activity* of a selected sustainability category, providing a focal point for an action to address or engage with. Examples include, “solar panels” (entity of *resource optimization*), “carbon reduction targets” (goal of *emissions control*), “wildlife conservation” (sub-area of *Ecological Conservation*), “workplace safety sessions” (activity of *Worker & Consumer Safety*).
3. We provide a list of sustainability categories that an aspect can belong to. These include the categories of *Resource Optimization*, *Waste*

*Management, Emissions Control, Ecological Conservation, Workplace, Outreach, Management, Business Compliance, Worker & Consumer Safety, Data & Cybersecurity Protection.* [Note: We provided the annotators with Table 5, which provides a full explanation of all the sustainability categories.]

4. The aspect must directly and explicitly relate to a specific sustainability category, requiring no additional assumptions to establish its relevance.
5. Where an aspect comprises multiple words, it must be a continuous sequence of words within the sustainability statement. Therefore, a single aspect cannot appear in two or more different parts of the sentence.
6. As much as possible, an aspect should be described with greater specificity, including additional terms such as qualifiers. An example of including qualifiers is the annotation of “workplace safety measures” instead of “safety measures”.
7. However, to maintain brevity, an aspect should avoid redundant words that do not add to its core meaning. For example, “safety protocols” should be annotated instead of “safety protocols in the event of similar incidents”.

#### **Guidelines for Action Annotation:**

1. Action labels characterize the type of action that is taken toward the aspect based on the sustainability statement. For each aspect, these action labels constitute one of the following: *planning, implemented, indeterminate*.
2. **Planning:** Indicates that an action has been planned or a commitment has been made by a company to address or engage with the aspect to advance its sustainability efforts. This can include:
  - Future plans to incorporate a sustainability aspect within sustainability operations, particularly in cases where the aspect is an activity or entity. For example, in the statement, “the company plans to increase waste recycling by 50% in 2025”, “waste recycling” is the aspect, and “planning” is its associated action.

- Future commitment with respect to a sustainability aspect, particularly in cases where the aspect is a goal or sub-area. For example, in the statement, “the group is firmly committed to improve the local community welfare this year”, “local community welfare” is the aspect, and “planning” is its associated action.

3. **Implemented:** Indicates that an action has already been taken to address or engage with the aspect to advance its sustainability efforts. This can include:

- A sustainability aspect has already been incorporated within sustainability operations, particularly in cases where the aspect is an activity or entity. For example, in the statement, “the company has successfully deployed wastewater filtration technologies”, “wastewater filtration technologies” is the aspect, and “implemented” is its associated action.
- A sustainability aspect has been advanced or achieved, particularly in cases where the aspect is a goal or sub-area. For example, in the statement, “we have successfully achieved 20% lower carbon emissions in 2024”, “carbon emissions” is the aspect, and “implemented” is its associated action.

4. **Indeterminate:** Indicates that it is unclear from the statement if the company intends to address or engage with the aspect to advance its sustainability efforts, or how it intends to do so. This can include:

- Statements without clear attribution as to who is engaging with the aspect, casting doubt over whether the company is explicitly involved in the engagement. For example, in the statement “Investments are being made in the country’s renewable energy projects”, “renewable energy projects” is the aspect, while “indeterminate” is its associated action.
- Statements that reflect uncertain or non-committal language with respect to engaging with the aspect, lacking the clarity required to indicate concrete commitments, plans or implementations. These include statements that possess hedging language that express uncertainty, without assurance

of follow-through. For example, the hedging phrase, “if feasible”, in the statement, “we intend to improve workplace diversity, if feasible”. For this statement, “workplace diversity” is the aspect, while “indeterminate” is its associated action. Other examples include the use of exploratory language, again, without any assurance of follow-through, such as “we are considering” in the statement “we are considering increasing the funding for our community development program”. For this statement, “community development program” is the aspect, while “indeterminate” is its associated action.

- Statements that do not highlight an enhancement in sustainability. These can include mere disclosures of figures, such as, “we will disclose greenhouse emission numbers from next year”, where “greenhouse emission numbers” is the aspect, and “indeterminate” is its associated action. Additionally, simply listing responsibilities without an explicit link to improvements in sustainability. For example, “the board is responsible for corporate governance”, where “corporate governance” is the aspect, and “indeterminate” is its associated action.

## B Sustainability Categories

Table 5 shows the sustainability categories of the A3CG dataset, as well as their corresponding definitions. These categories align with standard sustainability taxonomies such as SASB (Eng et al., 2022) to remain relevant to corporate sustainability stakeholders.

## C Experiments

### C.1 Full Details on Data Partitions

Table 6 highlights the categories that constitute the different data splits discussed in section 4.1, and specifies the number of samples for each data partition. For each fold, seen categories correspond to the categories found in the seen train, seen val, seen test sets, while unseen categories correspond to the categories found in the unseen test set. Notably, the seen and unseen categories vary across folds to ensure a balanced evaluation. Additionally, the samples are split to ensure a train-validation-test ratio of roughly 68:11:21, with the test set comprising either seen or unseen categories.

We further breakdown the different aspect counts corresponding to the categories in Table 7. For each partition of train, validation, test set, we ensure that no aspect samples constitute roughly 22% of the total count of aspect + no aspects. Consistent with popular aspect-based analysis SemEval datasets (Pontiki et al., 2014, 2015, 2016), the inclusion of no aspect samples simulates real-world conditions, whereby some sustainability statements may not contain any aspects (Khan et al., 2016). Therefore, models encounter the realistic challenge of distinguishing sustainability statements that contain aspect-action pairs from those that do not, improving the robustness of our study.

Models are trained and validated exclusively on seen categories (i.e. the seen train and seen validation sets) only, and are then tested on unseen test sets. Further evaluation on the seen test set serves as a control experiment. This setup effectively prevents data leakage from the unseen categories during testing, ensuring the robustness of the results for unseen category generalization.

### C.2 Model Implementation Details

To adapt the SOTA-ABSA methods - InstructABSA (Scaria et al., 2023), IT-RER-ABSA (Zheng et al., 2024), GRACE (Luo et al., 2020), CONTRASTE (Mukherjee et al., 2023), the prediction target is changed to that of the aspect-action (aspect, action) pair in place of (aspect, sentiment). For CONTRASTE, the opinion terms are omitted in the prediction target, as well as the Opinion Term Detection module. For InstructABSA and IT-RER-ABSA, we leverage the original format of the prompt for the Sentiment-Pair Extraction (ASPE), adapting the prompt for aspect-action classification. All the prompts utilized in supervised models (InstructABSA, IT-RER-ABSA, T5, T5+AL, T5+CL,) can be found in Table 8, while the corresponding few-shot examples for InstructABSA can be found in Table 9. Few-shot examples for IT-RER-ABSA are retrieved from the training sets through the retrieval module in its original implementation. For vanilla models (T5, BERT-ST), as well as their implementation with learning paradigms (T5+AL, T5+CL, BERT-ST+AL, BERT-ST+CL), we use the T5-base<sup>5</sup> and BERT-base-uncased<sup>6</sup> pre-trained models from HuggingFace<sup>7</sup>. The implementation

<sup>5</sup><https://huggingface.co/google-t5/t5-base>

<sup>6</sup><https://huggingface.co/google-bert/bert-base-uncased>

<sup>7</sup><https://huggingface.co>




Figure 6: (Version with legend) t-SNE visualization of feature embeddings for fold 3 unseen samples, with each sample color-coded by category.

of our sequence tagging framework in BERT-ST, BERT-ST+CL, BERT-ST+AL, follows the joint BIO scheme in (Mao and Li, 2021), and is summarized in Table 3, which shows the extraction of the aspect-action pair (carbon emissions, implemented), from the input “our carbon emissions have reduced.”

Input	our	carbon	emissions	have	reduced
Aspect	O	B	I	O	O
Action	O	IMP	IMP	O	O

Table 3: Sequence Tagging using joint-BIO scheme. This example highlights the prediction of the aspect-action pair (carbon emissions, implemented).

For all LLMs utilized - GPT-4o (OpenAI, 2024), Claude 3.5 Sonnet (Huang et al., 2024b), Llama 3 (70B) (AI@Meta, 2024), and DeepSeek V3 (Liu et al., 2024), we leverage the prompt template in Table 8. We provide few-shot examples in Table 10 for the LLM few-shot setups. Few-shot examples are selected from the same categories as the test set but are not drawn from the test set itself. For consistency, the few-shot examples are kept the same across all LLMs tested. The versions of LLMs utilized include: gpt-4o-2024-08-06 for GPT-4o, claude-3-5-sonnet-20241022 for Claude 3.5 Sonnet, llama3-70b-8192 for Llama 3 (70B), DeepSeek-V3 2024/12/26 for DeepSeek V3.

### C.3 Hyperparameters for Supervised Models

To ensure the fairness and robustness of results (Xu et al., 2023b), we adopt the hyperparameter settings specified in the original papers for the ABSA-SOTA models - GRACE (Luo et al., 2020), IT-RER-ABSA (Zheng et al., 2024), InstructABSA (Scaria et al., 2023), CONTRASTE (Mukherjee et al., 2023). This ensures that the models reflect their intended configurations and design choices. For

T5 and BERT-ST, as well as their respective learning paradigm implementations (T5+AL, T5+CL, BERT-ST+AL, BERT-ST+CL), a learning rate of  $3e^{-5}$  was utilized, with a batch size of 8. For T5+CL and BERT-ST+CL, the models were pre-trained using contrastive learning, with a temperature,  $\tau$ , of 0.07, a maximum number of positives and negatives of 2 each, and pre-training over 20 epochs with early stopping (patience of 3). For adversarial training setups (T5+AL, BERT-ST+AL), the model might be highly sensitive to the parameter  $\alpha$ , which denotes the gradient scaling factor used in the gradient reversal layer (i.e. the strength of adversarial signal). For adversarial setups, we utilize a reasonable default  $\alpha$  of 0.3, while conducting a hyperparameter study on  $\alpha$  in appendix C.4. All models are trained over 100 epochs on the training set, where the model instance with the best score on the validation set is used for test set evaluation.

### C.4 Ablation Study for Adversarial Setups

$$\nabla_{\theta} \mathcal{L} = \nabla_{\theta} \mathcal{L}_{\text{task}} - \alpha \nabla_{\theta} \mathcal{L}_{\text{adv}} \quad (1)$$

In this study, the implementation of adversarial learning methods (T5+AL, BERT-ST+AL) follows from (Ganin and Lempitsky, 2015), where the objective is for the encoder to learn category-invariant features. First, a category discriminator is added to the encoder to predict the categories present in a given sample. Then, a gradient reversal layer is added to reverse the gradients resultant from the discriminator loss. We summarize this in equation (1) where  $L_{\text{adv}}$  represents the discriminator loss, and  $\alpha$  represents the gradient scaling factor which controls the strength of the adversarial signal. The total loss  $L$  combines the task loss  $L_{\text{task}}$  with a scaled  $L_{\text{adv}}$ , to ensure that the encoder learns category-invariant representations while balancing the learn-



ing of the primary task. Following from (Ganin and Lempitsky, 2015),  $\alpha$  controls the trade-off between the primary task objective and the adversarial objective (learning category-invariant features).

Since the main results of this study focus on the performance on unseen categories, tuning  $\alpha$  on the validation set (comprising seen categories) would not necessarily optimize performance on the unseen test set (comprising unseen categories). Therefore, we selected a default  $\alpha$  value of 0.3 for all adversarial setups (T5+AL, BERT-ST+AL). To validate the robustness of this choice, we observe how varying  $\alpha$  impacts F1 for BERT-ST and T5 on the unseen test set. We examine  $\alpha \in \{0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 1.9, 2.0\}$ , a range that broadly aligns with foundational adversarial learning setups (Ganin et al., 2016).




Figure 7: F1 results averaged across unseen test sets for all folds (US Avg), for T5+AL and BERT-ST+AL, corresponding to different values of  $\alpha$ .

From Figure 7, we observe that even in an ideal scenario where the chosen  $\alpha$  value maximizes performance across all unseen test sets, the US Avg (average f1 across unseen (US) test sets of all folds) only reaches a maximum value of 43.19 for T5+AL ( $\alpha = 0.6$ ) and 27.47 for BERT-ST+AL ( $\alpha = 0.5$ ). However, note that this represents only the ideal  $\alpha$  value, which may not be achievable in practice without tuning on the unseen test sets - an approach that would constitute an unfair implementation. Moreover, even with these ideal  $\alpha$  values, the observed gains over the experimentally selected  $\alpha$  value of 0.3 - which yielded 42.82 for T5+AL and 26.24 for BERT-ST+AL (Table 2), remain marginal and are still lower than the results achieved with T5+CL and BERT-ST+CL. As a result, the ablation study underscores that while adversarial learning

still provides benefits for cross-category generalization in A3CG, contrastive learning is relatively more effective. This aligns with our discussion in section 5.2. We provide the full breakdown of US Avg F1 results for different  $\alpha$  in Table 11.

## C.5 Error Analysis Examples

All error examples discussed in the main text are outlined in tables 12 and 13

## C.6 Results Breakdown

The detailed results breakdown for all methods corresponding to each subtask of A3CG is found in Table 14.

## D Computation and Tools Used for Our Study

This study was conducted with the help of external, publicly available tools (Pytorch, Huggingface, GPT-4o, Llama3, DeepSeek V3, Claude 3.5 Sonnet), with all experiments run on a single NVIDIA GeForce RTX 4090 GPU. We use the T5-base<sup>8</sup> version (for T5, T5+AL, T5+CL) and BERT-base-uncased<sup>9</sup> version (for BERT-ST, BERT-ST+AL, BERT-ST+CL) from Huggingface<sup>10</sup>. GPT-4o is deployed through the OpenAI API<sup>11</sup>, Llama3 (70B) through the Groq API<sup>12</sup>, DeepSeek V3 through the DeepSeek API<sup>13</sup>, Claude 3.5 Sonnet through the Anthropic API<sup>14</sup>. We specify the model sizes in the following (model sizes for proprietary models are unavailable): BERT(110M), T5(220M), Llama3 (70B), DeepSeek V3 (671B), while the model sizes for SOTA-ABSA methods can be found from their respective papers - GRACE (Luo et al., 2020), IT-RER-ABSA (Zheng et al., 2024), InstructABSA (Scaria et al., 2023), CONTRASTE (Mukherjee et al., 2023).

<sup>8</sup><https://huggingface.co/google-t5/t5-base>

<sup>9</sup><https://huggingface.co/google-bert/bert-base-uncased>

<sup>10</sup><https://huggingface.co>

<sup>11</sup><https://openai.com/index/openai-api/>

<sup>12</sup><https://console.groq.com/docs/quickstart>

<sup>13</sup><https://platform.deepseek.com/>

<sup>14</sup><https://www.anthropic.com/api>

Sustainability Statement	Labels
We strive to improve patient safety through adequate trainings for our people as well as maintaining open communications with our patients.	("patient safety", "planning") - aspect category: <b>Worker &amp; Consumer Safety</b> ; ("adequate trainings", "planning") - aspect category: <b>Workplace</b>
We have also embarked on conservation efforts outside of our oil palm plantations.	("conservation efforts", "implemented") - aspect category: <b>Ecological Conservation</b>
The carbon footprint intensity decreased 5% from the previous year mainly due to continuous efforts to optimize eco-efficiency in our building operations.	("carbon footprint intensity", "implemented") - aspect category: <b>Emissions Control</b> ; ("eco-efficiency", "implemented") - aspect category: <b>Resource Optimization</b>
It is imperative that we manage our business prudently with high standard of corporate governance and integrity.	("corporate governance", "indeterminate") - aspect category: <b>Management</b>
The Group aims to achieve this by implementing several ethics and anti-corruption policies across our operations.	("anti-corruption policies", "planning") - aspect category: <b>Business Compliance</b>
With a PDPA officer appointed, we continue to maintain high standards in safekeeping of personal data of our patients.	("personal data", "implemented") - aspect category: <b>Data &amp; Cybersecurity Protection</b> ;
Factory closures, as well as labour and capacity shortages have resulted in volatile freight costs and uncertain delivery times.	<b>no aspect &amp; no action</b>

Table 4: Samples from the A3CG dataset. Outputs are in the format of (**aspect, action**) pair(s), and the category of the aspect is also denoted. Pairs, each comprising of an aspect-action pair and aspect category, are delineated by ;

<b>Resource Optimization</b>	Prioritizes the proactive, efficient, and sustainable selection and use of energy, water, and materials. This approach focuses on minimizing environmental impact through the strategic choice and utilization of resources. Distinct from waste management as it does not primarily deal with waste but rather prevents waste generation by optimizing resource use from the start.
<b>Waste Management</b>	Concentrates on managing waste after it has been produced, including the reduction, handling, disposal, and treatment of solid waste and wastewater. It involves but is not limited to recycling, reusing materials, and employing wastewater treatment processes to mitigate environmental degradation and facilitate resource recovery.
<b>Emissions Control</b>	Targets the reduction of various emissions from company operations, with an emphasis on mitigating climate change impacts or protecting air quality.
<b>Ecological Conservation</b>	Manages the company's interaction with natural ecosystems, focusing on minimizing harmful impacts and promoting biodiversity and habitat conservation.
<b>Workplace</b>	Focuses on the social aspects of sustainability within the company to promote equity, inclusivity, or quality of life, emphasizing employee well-being, inclusivity, or professional growth.
<b>Outreach</b>	Enhances social sustainability by contributing to community well-being, fostering customer satisfaction, building customer relationships, or collaborating with strategic partners to promote equity, inclusivity, and quality of life.
<b>Management</b>	Ensures strategic, ethical, or effective governance through rigorous oversight, diverse board composition, active stakeholder engagement, or high ethical standards within company leadership.
<b>Business Compliance</b>	Adherence to anti-corruption, anti-competitive practices, financial regulations, or financial reporting.
<b>Worker &amp; Consumer Safety</b>	Adherence to laws that guarantee the safety of products and services to consumer health, as well as fair labor practices and the protection of workers' rights.
<b>Data &amp; Cybersecurity Protection</b>	Protection and confidentiality of data, which includes measures against cyber threats and compliance with data security regulations.

Table 5: Sustainability Categories

Data Split	Full Dataset	Fold 1	Fold 2	Fold 3
Seen Categories	All categories	Worker & Consumer Safety, Workplace, Outreach, Management, Business Compliance, Data And Cybersecurity Protection, Ecological Conservation	Emissions Control, Resource Optimization, Outreach, Management, Business Compliance, Waste Management	Resource Optimization, Worker & Consumer Safety, Emissions Control, Workplace, Waste Management, Data And Cybersecurity Protection, Ecological Conservation
Unseen Categories	-	Emissions Control, Waste Management, Resource Optimization	Worker & Consumer Safety, Data And Cybersecurity Protection, Workplace, Ecological Conservation	Outreach, Management, Business Compliance
Seen Train Samples	1391	850	850	850
Seen Val Samples	214	135	135	135
Seen Test Samples	399	270	270	270
Unseen Test Samples	-	270	270	270
Total Samples	2004	1525	1525	1525

Table 6: Detailed Breakdown of Data Splits. Seen Categories correspond to the categories found within Seen Train, Seen Validation, Seen Test Samples. Unseen Categories correspond to the categories found within Unseen Test Samples. Seen Train, Seen Validation, Seen Test Samples, Total Samples and Unseen Test Samples correspond to the number of samples (i.e. statements) for each data split.

Data Split/ Aspect Category	Full Dataset			Fold 1				Fold 2				Fold 3			
	Seen Train	Seen Val	Seen Test	Seen Train	Seen Val	Seen Test	Unseen Test	Seen Train	Seen Val	Seen Test	Unseen Test	Seen Train	Seen Val	Seen Test	Unseen Test
Resource Optimisation	266	33	69	-	-	-	126	199	45	86	-	238	35	62	-
Management	209	25	63	187	26	62	-	154	29	50	-	-	-	-	101
Workplace	277	43	81	264	50	66	-	-	-	-	164	248	19	71	-
Emissions Control	200	34	50	-	-	-	99	175	28	51	-	165	42	47	-
Worker & Consumer Safety	222	35	71	225	28	65	-	-	-	-	131	193	30	57	-
Business Compliance	112	24	31	126	10	29	-	79	16	37	-	-	-	-	67
Outreach	234	35	71	200	38	73	-	201	23	54	-	-	-	-	127
Waste Management	212	29	66	-	-	-	102	187	36	56	-	180	23	63	-
Data & Cybersecurity Protection	93	16	20	92	13	24	-	-	-	-	58	81	12	13	-
Ecological Conservation	64	10	28	64	5	17	-	-	-	-	27	50	11	16	-
Total No Aspect	463	73	130	307	47	94	94	297	47	94	94	297	47	94	94
Total Aspect Count*	1889	284	550	1158	170	336	327	995	177	334	380	1155	172	329	295

Table 7: Aspect counts for each of the different aspect categories within each split of the dataset. \*The total aspect count excludes the count of no aspects. For each data split, the number of no aspect samples are kept at a consistent percentage - around 22% of the sum of total number of aspects + no aspects.

Method(s)	Prompt
T5, T5+AL, T5+CL	Extract aspect-action pairs from the following sentence: { <b>Sustainability Statement</b> }
InstructABSA, IT-RER-ABSA	<p>Definition: The output will be the aspects and the aspect’s action classification. In cases where there are no aspects the output should be noaspectterm:none.</p> <p>Positive example 1- {<b>Few-Shot Example</b>}</p> <p>Positive example 2- {<b>Few-Shot Example</b>}</p> <p>Now complete the following example- Input: {<b>Sustainability Statement</b>}</p>
All LLMs (GPT-4o, Claude 3.5 Sonnet, Llama 3 (70B), DeepSeek V3)	<p>[Task Description] The aspect-action pair consists of aspect term and action label. You are tasked with extracting aspect-action pairs from sentences.</p> <p>[Definitions of Aspects] Aspects refer to sustainability-related elements found within the sentence that describe entities, goals, sub-areas, activities of sustainability. To enhance specificity, aspects should incorporate appropriate qualifiers that clarify their meaning. At the same time, they should remain succinct by eliminating any redundant language that do not contribute to their core meaning. Aspects must be explicitly found within the sentence and cannot include words not present in the sentence.</p> <p>[Definitions of Actions] Action labels are the type of action taken toward the aspect. Action labels fall under the following categories: “Planning”: Indicates that an action has been planned or a commitment has been made by a company to address or engage with the aspect to advance its sustainability efforts. “Implemented”: Indicates that an action has already been taken to address or engage with the aspect to advance the company’s sustainability efforts. “Indeterminate”: Indicates that it is unclear from the statement if the company intends to address or engage with the aspect to advance its sustainability efforts, or how it intends to do so.</p> <p>[Response Instructions] Evaluate the explicit content of the sentence to determine the aspect-action pairs. There can be more than one aspect-action pair the sentence. However, each aspect can only have one action. Alternatively, the sentence may not have any aspect-action pairs. In such cases, return (“no aspect”, “no action”).</p> <p>[Response format] For the sentence, provide the aspect-action pairs in tuples using the following format. Do not output any explanations. You can observe the following example(s) for reference.</p> <p>{<b>Few-Shot Examples</b>}</p> <p>Now complete the following example- Input: {<b>Sustainability Statement</b>}</p>

Table 8: Prompt templates for all methods that require prompting. “Sustainability statement” refers to the dataset samples. For InstructABSA & IT-RER-ABSA, the prompt is kept similar to the original format of the Sentiment-Pair Extraction (ASPE) task in the respective papers. Few-shot examples for InstructABSA is found in Table 9. Few-shot examples for LLMs are found in Table 10. Few-shot examples for IT-RER-ABSA are retrieved from the training sets through the retrieval module in its original implementation.



<b>Dataset Partition</b>	<b>Examples</b>
Full Dataset	<p>Input: For example, in the case of Flu-Pandemic disruption, the SaFe Management Measures System will be activated to prevent pandemic virus transmission at the workplace. Output: pandemic virus transmission:planning, workplace:planning</p> <p>Input: We encourage our team to reduce paper usage and to reuse or recycle non-sensitive paper waste where practicable. Output: paper usage:indeterminate, non-sensitive paper waste:indeterminate</p>
Fold 1	<p>Input: For example, in the case of Flu-Pandemic disruption, the SaFe Management Measures System will be activated to prevent pandemic virus transmission at the workplace. Output: pandemic virus transmission:planning, workplace:planning</p> <p>Input: Besides scouting for the latest technology and solutions, the company also engages in strategic partnerships with research institutions. Output: strategic partnerships:implemented</p>
Fold 2	<p>Input: We reviewed our regulatory risks as part of the management risk assessment process in 2021 and have achieved the 2021 target. Output: management risk assessment process:implemented, regulatory risks:implemented</p> <p>Input: We strive to adopt best practices to optimize consumption and minimise carbon emissions as an organisation. Output: carbon emissions:planning</p>
Fold 3	<p>Input: We are committed to improve on our occupational health and safety initiatives and conduct regular reviews of our programmes, processes, risk assessments and controls. Output: occupational health and safety initiatives:planning, risk assessments and controls:planning</p> <p>Input: We encourage our team to reduce paper usage and to reuse or recycle non-sensitive paper waste where practicable. Output: paper usage:indeterminate, non-sensitive paper waste:indeterminate</p>

Table 9: Few-Shot Examples for InstructABSA

Data Partition	Example
Full Dataset	<p>Input: The Group revitalised and rejuvenated five existing coffeeshops in FY2022 to improve customers' dining experience and hygiene standards. Response: ("hygiene standards", "implemented"), ("customers' dining experience", "implemented")</p> <p>Input: It is imperative that we manage our business prudently with high standard of corporate governance and integrity. Response: ("corporate governance", "indeterminate")</p> <p>Input: We are committed to improve on our occupational health and safety initiatives and conduct regular reviews of our programmes, processes, risk assessments and controls. Response: ("occupational health and safety initiatives", "planning"), ("risk assessments and controls", "planning")</p> <p>Input: This project will commence in 2022 and take place over the next 3 years. Response: ("no aspect", "no action")</p> <p>Input: In 2020, we achieved a 44% reduction in carbon emissions intensity against 2007 levels, putting us on track to achieving our SBTi-validated target of 59% by 2030. Response: ("carbon emissions intensity", "implemented")</p>
Fold 1 Seen Test	<p>Input: This project will commence in 2022 and take place over the next 3 years. Response: ("no aspect", "no action")</p> <p>Input: We have in place robust worksite inspection procedures and monthly audits to identify workplace hazards and ensure all activities comply with all Group and regulatory requirements. Response: ("workplace hazards", "implemented"), ("worksite inspection procedures", "implemented")</p> <p>Input: Moving ahead, we are targeting to include automation in tracking order status and completion, routing and invoicing to improve customer experience. Response: ("customer experience", "planning")</p> <p>Input: Continuous learning is necessary to help enhance workmen proficiency. Response: ("workmen proficiency", "indeterminate"), ("continuous learning", "indeterminate")</p> <p>Input: Biodiversity conservation efforts shaping up across plants. Response: ("biodiversity conservation efforts", "indeterminate")</p>
Fold 1 Unseen Test	<p>Input: The upgrading programme of our second generation fleet of 19 C651 trains has commenced and is expected to be completed in 2018. Response: ("no aspect", "no action")</p> <p>Input: In 2020, we achieved a 44% reduction in carbon emissions intensity against 2007 levels, putting us on track to achieving our SBTi-validated target of 59% by 2030. Response: ("carbon emissions intensity", "implemented")</p> <p>Input: The carbon footprint intensity decreased 5% from the previous year mainly due to continuous efforts to optimize eco-efficiency in our building operations. Response: ("carbon footprint intensity", "implemented"), ("eco-efficiency", "implemented")</p> <p>Input: In 2019, the Group generated a total of approximately 936,000 tonnes of sludge. Response: ("sludge", "indeterminate")</p> <p>Input: We will continue to maintain or lower our energy and water consumption in FY2022. Response: ("energy and water consumption", "planning")</p>

Table 10: Few-Shot Examples for LLMs - continued on the next page

Data Partition	Example
Fold 2 Seen Test	<p>Input: Each goal has specific targets to be achieved over the next 15 years. Response: ("no aspect", "no action")</p> <p>Input: In FY2019, apart from economic performance, we have with the help of an independent external consultant established our sustainability performance management framework. Response: ("sustainability performance management framework", "implemented")</p> <p>Input: We remain committed in ensuring that we handle medical waste in a safe and sustainable manner. Response: ("medical waste", "planning")</p> <p>Input: In 2019, the Group generated a total of approximately 936,000 tonnes of sludge. Response: ("sludge", "indeterminate")</p> <p>Input: Free shuttle bus services are also available at some of its shopping malls. Response: ("no aspect", "no action")</p>
Fold 2 Unseen Test	<p>Input: We have in place robust worksite inspection procedures and monthly audits to identify workplace hazards and ensure all activities comply with all Group and regulatory requirements. Response: ("workplace hazards", "implemented"), ("worksite inspection procedures", "implemented")</p> <p>Input: This project will commence in 2022 and take place over the next 3 years. Response: ("no aspect", "no action")</p> <p>Input: Continuous learning is necessary to help enhance workmen proficiency. Response: ("workmen proficiency", "indeterminate"), ("continuous learning", "indeterminate")</p> <p>Input: We have also embarked on conservation efforts outside of our oil palm plantations. Response: ("conservation efforts", "implemented")</p> <p>Input: We are committed to foster a non-discriminatory workplace environment. Response: ("non-discriminatory workplace environment", "planning")</p>
Fold 3 Seen Test	<p>Input: In 2020, we achieved a 44% reduction in carbon emissions intensity against 2007 levels, putting us on track to achieving our SBTi-validated target of 59% by 2030. Response: ("carbon emissions intensity", "implemented")</p> <p>Input: We are committed to improve on our occupational health and safety initiatives and conduct regular reviews of our programmes, processes, risk assessments and controls. Response: ("occupational health and safety initiatives", "planning"), ("risk assessments and controls", "planning")</p> <p>Input: We encourage our team to reduce paper usage and to reuse or recycle non-sensitive paper waste where practicable. Response: ("paper usage", "indeterminate"), ("non-sensitive paper waste", "indeterminate")</p> <p>Input: In order to achieve carbon reduction goals, the governments and relevant regulatory agencies have introduced successively various policies and measures to encourage various industries in implementing green and low-carbon transformation. Response: ("carbon reduction goals", "indeterminate"), ("green and low-carbon transformation", "indeterminate")</p> <p>Input: It covers the period from 1 May 2020 to 30 April 2021. Response: ("no aspect", "no action")</p>

Table 10: Few-Shot Examples for LLMs - continued on the next page

Data Partition	Example
Fold 3 Unseen Test	Input: Each goal has specific targets to be achieved over the next 15 years. Response: ("no aspect", "no action")
	Input: It is imperative that we manage our business prudently with high standards of corporate governance and integrity. Response: ("corporate governance", "indeterminate")
	Input: In Singapore, our community programmes and approach are aligned with NCSS' Strategic Thrusts where every person is empowered to live with dignity in a caring and inclusive society. Response: ("community programmes", "implemented")
	Input: We have also established a business continuity plan (BCP) which focuses on the recovery of technology facilities and platforms, such as critical applications, databases, servers or other required technology infrastructure for the viability of the business. Response: ("business continuity plan", "implemented")
	Input: Ranked 2nd in customer satisfaction for four consecutive years among supermarket retailers in Singapore. Response: ("customer satisfaction", "indeterminate")

Table 10: Few-Shot Examples for LLMs

$\alpha$	T5+AL				BERT-ST+AL			
	Fold 1 US	Fold 2 US	Fold 3 US	US Avg	Fold 1 US	Fold 2 US	Fold 3 US	US Avg
0.01	45.30	44.06	39.42	42.93	26.31	21.3	29.99	25.87
0.05	42.86	46.3	41.67	43.61	24.09	21.9	31.2	25.73
0.1	40.27	44.82	38.77	41.29	23.28	21.49	29.31	24.69
0.2	36.56	47.78	41.12	41.82	26.17	21.82	28.78	25.59
0.3	39.62	47.02	41.82	<u>42.82</u>	27.05	25.57	26.11	<u>26.24</u>
0.4	40.88	46.27	41.05	42.73	22.63	24.12	29.31	25.35
0.5	40.78	43.85	34.51	39.71	27.53	25.57	29.31	<b>27.47</b>
0.6	41.63	45.05	42.90	<b>43.19</b>	27.53	20.65	29.31	25.83
0.7	42.94	44.77	35.63	41.11	27.53	25.57	29.31	<b>27.47</b>
0.8	36.08	45.54	37.98	39.87	27.53	22.52	29.31	26.45
0.9	42.05	44.9	35.22	40.72	27.53	21.04	29.31	25.96
1.0	37.45	43.39	37.39	39.41	27.53	25.57	29.31	<b>27.47</b>
1.1	39.39	44.69	38.7	40.93	22.63	25.57	31.74	26.65
1.2	42.8	45.78	35.88	41.49	27.53	25.57	29.31	<b>27.47</b>
1.3	40.0	45.21	36.58	40.60	27.53	25.57	29.31	<b>27.47</b>
1.4	37.36	45.88	36.68	39.97	27.53	25.57	29.31	<b>27.47</b>
1.5	42.18	45.32	40.4	42.63	27.53	25.57	28.06	27.05
1.6	41.85	45.34	41.69	42.96	27.53	25.57	29.31	<b>27.47</b>
1.7	36.08	45.22	32.93	38.08	27.53	25.57	29.31	<b>27.47</b>
1.8	36.89	43.14	38.34	39.46	27.53	20.72	29.31	25.85
1.9	35.13	43.14	37.37	38.55	27.53	25.57	29.31	<b>27.47</b>
2.0	40.87	44.14	35.56	40.19	27.53	25.57	29.31	<b>27.47</b>

Table 11: F1 performance of T5+AL and BERT-ST+AL across different  $\alpha$  values, for unseen (US) test sets of different folds. US Avg represents average across unseen (US) test sets of all folds. Best US Avg results across all  $\alpha$  values are bolded. Results of experimentally selected  $\alpha$ , 0.3, are underlined.



<b>Error Type</b>	<b>Sample</b>	<b>Analysis</b>
Missing Obvious Aspects	Our innovations also contribute to the reduction of carbon emissions from our delivery fleet due to fewer redeliveries required	When generalizing to an unseen sustainability category, supervised models cannot extract sustainability aspects (i.e. “carbon emissions”) that are obvious and highly intuitive to sustainability.
i) Elipsis	The company strictly abides by the energy conservation law of the PRC, strives to continuously improve resource efficiency and reduce resource consumption.	All supervised models fail to classify the action associated with the “reduce consumption” aspect as “planning”. This highlights the inability of supervised models to resolve ellipsis. Specifically, “strives to” applies to both “resource efficiency and “resource consumption”, but does not directly precede the latter, therefore requiring contextual inference.
ii) Ambiguous Syntax	During FY2022, we shared a total of S\$129,000 in direct economic value towards community investments, with the decrease in activities related to the pandemic, which restricted face-to-face community interaction.	All supervised models fail to classify the action associated with the “community investments” aspect as “indeterminate.” This highlights their inability to resolve syntactic ambiguity. Specifically, the statement contains an ambiguous modification: it is unclear whether the phrase “with the decrease in activities related to the pandemic” modifies “S\$129,000 in direct economic value towards community investments”, potentially implying a decrease in community investment. In such cases, where ambiguity is present, models fail to classify the action as “indeterminate”.

Table 12: Error Analysis for Supervised Models

<b>Error Type</b>	<b>Sample</b>	<b>Analysis</b>
i) Misinterpreting Modality	Where possible, we have implemented sustainable measures to monitor our water consumption and increase water efficiency.	LLMs fail to classify “indeterminate” as the action associated with the aspects of “water consumption” and “water efficiency” respectively. LLMs fail to recognise that this statement represents uncertainty and tentativeness, misinterpreting the modality involved. Specifically, “where possible” depicts hedging and non-committal language.
ii) Negation Handling	The additional safety measures previously implemented in 2021 also have been reduced or removed accordingly based on each country’s local laws and regulations.	LLMs fail to recognise that the “safety measures” have been reduced or removed, demonstrating no enhancement in safety measures, and thereby no improvement in sustainability efforts. By failing to handle this negative case, LLMs fail to classify “indeterminate” as the action associate with the aspect “safety measures”.
iii) Unattributed statements	It maps out a framework that a company or organization can follow to set up an effective environmental management system.	LLMs fail to classify “indeterminate” for the aspect “environmental management system”. LLMs fail to recognise that this statement is unattributed, meaning that it is unclear who is engaging with the aspect “environmental management system”. In statements such as these, it is unclear whether the company is directly engages with the aspect, warranting an “indeterminate” classification for the aspect.
iv) Future Dependency	When all of the company’s Green Mark-awarded projects are fully completed, they will contribute to a reduction in carbon emissions of almost 95,000 tonnes annually.	LLMs fail to classify “planning” for the aspect “carbon emissions”. LLMs fail to recognize that reducing carbon emissions is a future-oriented goal that is contingent on a forthcoming event. The goal can only be achieved upon the full completion of the the company’s Green Mark-awarded projects, highlighting its forward-looking nature.

Table 13: Error Analysis for LLMs

Method	Metric	Full Dataset	Fold 1		Fold 2		Fold 3		S Avg	US Avg
			S	US	S	US	S	US		
T5	ATE	76.28	70.49	48.34	77.82	53.04	78.95	38.14	75.75	46.51
	AC	64.19	49.85	31.00	62.52	34.72	61.30	15.51	57.89	27.08
	AAA	70.48	57.85	43.03	68.90	45.74	67.94	34.59	64.90	41.12
BERT-ST	ATE	53.55	50.67	27.54	52.75	24.77	53.27	33.58	52.23	28.63
	AC	31.51	27.39	2.58	24.05	3.36	30.79	8.77	27.41	4.90
	AAA	43.19	39.56	25.25	39.00	22.92	43.97	30.01	40.84	26.06
T5 + CL	ATE	77.05	73.38	52.65	78.54	52.48	77.80	40.92	76.57	48.68
	AC	65.72	56.80	37.02	63.19	36.87	61.40	22.34	60.46	32.08
	AAA	71.12	62.96	46.97	69.76	46.67	67.99	38.33	66.90	43.99
T5 + AL	ATE	75.69	71.77	43.17	77.24	54.08	76.98	45.80	75.33	47.68
	AC	64.27	53.87	26.58	59.65	37.29	58.59	29.12	57.37	31.00
	AAA	69.27	61.24	39.62	66.91	47.02	65.17	41.82	64.44	42.82
BERT-ST + CL	ATE	79.40	74.82	39.32	79.67	52.82	78.11	37.95	77.53	43.36
	AC	62.41	51.77	21.05	63.60	30.76	48.20	12.98	54.52	21.60
	AAA	68.53	60.00	37.06	69.22	41.78	58.04	34.94	62.42	37.93
BERT-ST + AL	ATE	24.30	48.77	27.63	27.13	25.57	46.51	27.60	40.80	26.93
	AC	0.00	23.75	3.19	0.00	0.00	19.28	3.79	14.34	2.33
	AAA	24.30	37.75	27.05	27.13	25.57	34.70	26.11	33.19	26.24
InstructABSA	ATE	75.22	70.47	40.55	74.90	55.90	73.33	51.24	72.90	49.23
	AC	65.28	53.07	23.44	57.46	40.45	58.15	38.33	56.23	34.07
	AAA	69.47	60.23	37.53	64.73	49.76	64.14	47.38	63.03	44.89
IT-RER-ABSA	ATE	76.04	71.72	65.13	76.87	55.73	77.87	41.94	75.49	54.27
	AC	63.64	50.72	30.20	59.91	39.54	63.46	24.24	58.03	31.33
	AAA	69.20	57.70	41.81	66.02	48.87	68.83	39.10	64.18	43.26
GRACE	ATE	75.45	74.53	59.15	76.59	51.98	76.90	50.38	76.01	53.84
	AC	59.48	50.39	30.50	53.67	30.13	51.41	20.69	51.82	27.11
	AAA	67.09	60.87	50.08	63.10	44.33	61.92	48.12	61.96	47.51
CONTRASTE	ATE	78.32	75.43	54.55	78.02	60.07	80.61	42.03	78.02	52.22
	AC	65.14	54.28	40.69	64.61	42.45	66.46	26.46	61.78	36.53
	AAA	71.33	61.26	48.14	69.81	50.30	71.34	40.58	67.47	46.34
GPT-4o	ATE	42.30	40.90	52.59	49.09	41.44	46.23	39.10	45.41	44.38
	AC	26.12	24.25	37.52	32.95	25.08	30.42	20.93	29.21	27.84
	AAA	29.79	31.61	42.98	40.00	32.51	35.58	32.35	35.73	35.95
GPT-4o + FS	ATE	48.92	47.04	59.68	48.76	51.71	52.16	42.08	49.32	51.16
	AC	30.70	32.44	38.45	29.95	36.82	36.96	25.47	33.12	33.58
	AAA	35.69	39.08	46.68	39.12	41.10	40.05	33.46	39.42	40.41
Claude 3.5 Sonnet	ATE	48.90	46.50	53.13	52.96	46.24	51.82	44.14	50.43	47.84
	AC	34.57	31.38	34.36	36.87	33.13	39.14	32.01	35.80	33.17
	AAA	37.70	36.71	39.44	41.11	36.88	41.59	38.22	39.80	38.18
Claude 3.5 Sonnet + FS	ATE	53.59	49.22	59.90	55.83	47.81	56.06	46.27	53.70	51.33
	AC	37.70	36.21	40.00	36.00	38.57	42.94	34.03	38.38	37.53
	AAA	42.11	40.62	46.27	43.18	40.35	45.04	39.48	42.95	42.03
Llama 3 (70B)	ATE	33.79	28.79	41.88	35.78	27.96	35.42	27.54	33.33	32.46
	AC	22.90	20.74	27.61	26.45	20.34	25.72	20.91	24.30	22.95
	AAA	20.15	17.97	25.24	23.66	18.33	22.31	18.43	21.31	20.67
Llama 3 (70B) + FS	ATE	43.27	37.13	52.89	42.33	42.01	49.72	34.99	43.06	48.80
	AC	30.09	23.86	32.96	29.05	33.41	30.47	24.76	27.79	30.38
	AAA	29.82	25.03	38.11	33.49	33.30	33.15	26.65	30.56	32.69
DeepSeek V3	ATE	47.29	45.77	57.47	52.99	47.44	51.76	41.49	50.17	48.40
	AC	33.13	27.59	41.46	36.21	33.19	36.12	23.49	33.31	32.71
	AAA	40.16	38.06	48.18	45.27	40.23	42.46	34.84	41.96	41.08
DeepSeek V3 + FS	ATE	43.27	47.90	37.48	54.02	55.89	55.01	42.86	52.31	45.41
	AC	30.09	31.90	17.09	36.67	39.61	42.96	28.69	37.18	28.46
	AAA	35.63	38.85	30.16	44.18	44.51	46.62	34.43	43.22	36.37

Table 14: ATE, AC, and AAA F1 scores, across the full dataset and folds with seen (S) and unseen (US) categories.