

Better Handling Coreference Resolution in Aspect Level Sentiment Classification by Fine-Tuning Language Models

Dhruv Mullick

Dept. of Computing Science
University of Alberta
mullick@ualberta.ca

Bilal Ghanem

Dept. of Computing Science
University of Alberta
bilalghm@gmail.com

Alona Fyshe

Dept. of Computing Science
University of Alberta
alona@ualberta.ca

Abstract

Customer feedback is invaluable to companies as they refine their products. Monitoring customer feedback can be automated with Aspect Level Sentiment Classification (ALSC) which allows analyzing specific aspects of the products in reviews. Large Language Models (LLMs) are the heart of many state-of-the-art ALSC solutions, but they perform poorly in some scenarios requiring Coreference Resolution (CR). In this work, we propose a framework to improve an LLM’s performance on CR-containing reviews by fine-tuning on highly inferential tasks. We show that the performance improvement is likely attributed to the improved model CR ability. We release a new dataset¹ that focuses on CR in ALSC, and share code² for the experiments.

1 Introduction

To understand an end user’s perspective on a product, it is common to consider reviews on online platforms. A company can look for the customers’ perspective on a certain aspect of the product. For instance, a laptop company might look for reviews concerning "battery." Aspect Level Sentiment Classification (ALSC) analyzes reviews for sentiments of specific aspects, like the "battery" aspect in earlier example (Yan et al., 2021). ALSC is a sub-task of a wider body of work called Aspect Based Sentiment Analysis (ABSA) (Liu, 2012), which aims to extract aspects and their associated sentiments. State-of-the-art ALSC solutions often use Large Language Models (LLMs) (Zhang et al., 2022).

Reviews often use pronouns, which can make coreference resolution (CR) in LLMs necessary to infer the sentiment associated with the aspect. Hence, LLMs used for ALSC need strong CR ability, and can fail otherwise. For instance, the sen-

¹<https://borealisdata.ca/dataset.xhtml?persistentId=doi:10.5683/SP3/HSKJEY>

²<https://github.com/dhruvmullick/absa-cs>

Table 1: Cases where the T5 ALSC model fails due to its poor coreference resolution ability.

Sentence	Aspect	Sentiment Polarity	
		Predicted	Gold
He ate food at the restaurant, it was <u>deserted</u>	restaurant	neutral	negative
	food	negative	neutral
He ate food at the restaurant, it was <u>dark</u>	restaurant	neutral	negative
	food	negative	neutral

tence - "*He ate food at the restaurant, it was deserted.*" requires the LLM to understand that the definite pronoun "it" refers to the "restaurant" (antecedent), because of the context ("deserted"). Table 1 shows four examples where the state-of-the-art T5 ALSC model (Zhang et al., 2021) fails due to its poor CR ability. We find that ~15% of this T5 model’s errors are on cases requiring CR ability.

LLMs are also known to have performance and stability issues (Phang et al., 2018). To remedy these, instead of directly training on the task of interest (target task), it can be beneficial to first train on an auxiliary task (Pruksachatkun et al., 2020). Certain auxiliary tasks can contribute to both improved performance and stability of the target task (Phang et al., 2018). Using auxiliary training, our work shows a way to improve an LLM’s performance on English ALSC reviews requiring CR.

In our work, we: **a)** show that an LLM trained for ALSC makes more errors when evaluated only on reviews requiring CR ability, compared to when handling typical ALSC reviews (8.7% mean F1); **b)** demonstrate that our framework for handling CR-containing reviews can improve ALSC model’s CR ability (16% mean F1); **c)** show that this improved CR ability can improve ALSC performance for reviews requiring CR ability (5% mean F1). **d)** release annotated variants of existing datasets which can be used to benchmark a model’s ALSC performance on CR cases.

2 Experimental Setup

2.1 Data

Original ALSC Datasets We consider English ALSC datasets: SemEval Restaurant (Rest16) (Pontiki et al., 2016) and MAMS (Jiang et al., 2019), both of which contain reviews from a similar restaurant domain. Inspired by Yan et al. (2021), ALSC reviews are processed into an input format suitable for our LLM - "[sentence]. aspect: [aspect]". The ground truth output is "positive", "negative" or "neutral". For example, "\$20 for good sushi cannot be beaten. aspect: sushi" has the ground truth as "positive". We clean datasets as per Appendix C.

CR Cases We identify reviews in the Rest16 and MAMS datasets that contain definite pronouns, and henceforth call these sentences *Pronoun cases*.

Limiting ourselves to the ALSC task described above, we say that a review is a *CR case* if its sentiment requires proper coreference resolution for correct classification. Specifically, *the aspect should be an antecedent of a definite pronoun which is associated with a sentiment polarity*. For example, "He ate food at the restaurant, it was deserted." with aspect: "restaurant" is a CR case. Here, "restaurant" is the antecedent of "it" which is associated with "deserted" and has negative connotations. CR cases are manually selected from Pronoun cases.

ALSC-CR Dataset Our dataset is composed of the original ALSC datasets (Rest16 and MAMS). The testing, however, is done only using CR cases, and we use a combination of Pronoun and Non-Pronoun cases for validation and train sets. Table 2 presents the dataset composition. Better performance on the test dataset will indicate a superior ability to handle CR cases in ALSC.

The train, validation and test sets are of similar, but not identical, distributions. Due to the limited number of CR cases, it is not possible to have train and validation sets composed entirely of CR cases. More details can be found in Appendix D.

2.2 Auxiliary Tasks

We use highly inferential tasks for auxiliary training in our experiments as they generally provide higher improvements for various NLP target tasks (Pruksachatkun et al., 2020). We select two commonsense tasks - CommonGen (Lin et al., 2020) and CosmosQA (Huang et al., 2019), as commonsense reasoning helps with CR (Liu et al., 2017).

SQuAD (Rajpurkar et al., 2016) is selected because it is a non-commonsense question answering (QA) task. Its performance is contrasted with CosmosQA, checking if it is the QA or the commonsense ability which improves CR. Quora Question Prediction (Wang et al., 2018) (QQP) is selected as it benefits performance on the Stanford Sentiment Treebank (SST) task which is similar to ALSC (Wang et al., 2019). Even if auxiliary tasks aren't designed for CR, they can impart CR ability to the model. For the QA example - "Context: Alice can't come. She is old"; "Question: Who is old?", answer is "Alice". Answering this requires CR and teaches the model CR ability.

CommonGen is a generative commonsense task involving sentence generation from a list of concepts (train size = 67,389). It tests: 1) ability to construct grammatical sentences adhering to commonsense; 2) reasoning with unseen concept combinations. For example: input - "concepts = [dog, frisbee, catch]"; output - "A dog leaps to catch a frisbee."

CosmosQA is a QA task where answering questions requires commonsense (train size = 25,262). For each question, there are four options, and the model should output the correct option number.

SQuAD is an extractive QA task where the correct answer to the question is present exactly in the passage (train size = 87,599).

QQP task involves checking if two Quora questions are semantically equivalent. We cap the train size at 50,000 to match the other datasets.

3 Experiments and Results

We ran experiments for three purposes: **a)** to show there is drop in ALSC performance for reviews requiring CR ability; **b)** to show we can alleviate this performance drop by auxiliary fine-tuning; **c)** to provide additional evidence that change in performance on CR cases is due to improved CR ability.

Inspired by state-of-the-art performance in Zhang et al. (2021), we used the T5 LLM (Raffel et al., 2019). Our baseline model is a T5 trained on ALSC-CR, but not fine-tuned on auxiliary tasks.

The T5 model was trained in various settings using training prompts/input prefixes (Appendix F). Wording of prompts has limited impact on the outcome so we did not experiment with the wording (Raffel et al., 2019). Rather, we relied on prior work for task prompts (Lin et al., 2020; Lourie et al., 2021; Raffel et al., 2019). For ALSC and

Table 2: ALSC-CR composition. Note that CR cases are types of Pronoun cases.

Partition	Size	Dataset		CR Cases	Data Type	
		MAMS	Rest16		Pronoun Cases	Non-Pronoun Cases
Train	12,434	✓	✓	✓	✓	✓
Validation	889	✓	✓	✓	✓	✓
Test	346	✓	✓	✓	✗	✗

Definite Pronoun Resolution (DPR) (Rahman and Ng, 2012) (Sec. 3.3), we created prompts as we did not find examples in prior work (see Appendix F).

All experiments were run with at least 10 random seeds, and Yuen-Welch test was used for testing statistical significance.

3.1 Model Performance on ALSC Without Auxiliary Fine Tuning

To check LLM performance on CR cases, we evaluated the T5 model on regular ALSC data (ALSC-Regular), which does not consist solely of CR cases. ALSC-Regular and ALSC-CR are equal sized and have an identical proportion of Rest16 and MAMS. We also evaluated the T5 model on ALSC-CR, to get the model’s performance solely on CR cases.

By comparing T5 model’s performance on the two ALSC datasets, we show that unspecialized LLMs face a significant performance problem while handling reviews requiring CR ability. Results are shown in Table 3, where evaluation on ALSC-CR shows a drop in performance of $\sim 8.7\%$ mean F1, as well as an increase of 0.6 F1 standard deviation indicating a poorer model convergence.

Table 3: T5 model evaluated on ALSC datasets. Best score bolded. Performances on the datasets are statistically significantly different ($p\text{-value}=9.03e-05$).

Dataset	Mean F1 (\pm Std. Dev)
ALSC-Regular	79.71 (± 1.99)
ALSC-CR	71.07 (± 2.60)

3.2 Fine Tuning With Auxiliary Tasks

As a solution to poor performance on ALSC-CR (Section 3.1), we experimented with various auxiliary tasks mentioned in Section 2.2.

We trained T5 model on the auxiliary task first to incorporate auxiliary task knowledge. This model is then trained and evaluated on ALSC-CR, our target task. We experimented with different auxiliary dataset sizes as the size has little correlation with the target task performance (Wang et al., 2019).

The model’s performance on ALSC-CR with different auxiliary tasks is compared to baseline model’s ALSC-CR performance to see if auxiliary tasks were beneficial. Results are shown in Table 4. We find that the lower ALSC-CR performance (compared to ALSC-Regular) can be alleviated by auxiliary training with CommonGen and QQP, which lead to statistically significant improvements of $\sim 5\%$ mean F1. Auxiliary training with CosmosQA and SQuAD does not lead to statistically significant improvement in any case.

Prior work (Pruksachatkun et al., 2020) showed a general improvement in a model’s target task performance when fine-tuned with highly inferential tasks. Apart from being highly inferential, because CommonGen is a generative commonsense task, it is ideal for imparting commonsense knowledge to a generative LLM like T5. On the other hand, CosmosQA being a discriminative task is unlikely to impart as much commonsense knowledge into a generative system (Lin et al., 2020). As being highly inferential is helpful for target tasks, the SQuAD extractive QA task, would not result in as significant an improvement. When used for auxiliary training, QQP shows a high improvement in the SST target task (Wang et al., 2019) which involves similar sentiment analysis, explaining QQP’s improved performance on ALSC-CR.

While auxiliary training on DPR appears promising, its dataset (train size = 1500) is much smaller than for other tasks. For completeness we did train using DPR but the mean F1 = 72.77 was not statistically significantly different from the baseline.

Similar to Wang et al. (2019), we do not find correlation between auxiliary task size and target performance. This lack of correlation may be due to the fact that small datasets might not teach the task sufficiently (Raffel et al., 2019). On the other hand, large auxiliary datasets can cause catastrophic forgetting of the LLM’s original objective (Wang et al., 2019). This original objective is generally beneficial for target tasks. Despite this lack of correlation, we have demonstrated a framework for improving

Table 4: Mean F1 (\pm Std. Dev) performance on ALSC-CR on different fractions of aux dataset. * denotes statistically significant difference from baseline. Table’s best scores bolded, 2^{nd} best underlined.

Aux. Task	Aux. Dataset Fraction			
	0.1	0.2	0.5	1.0
Commongen	<u>75.72</u> (\pm 1.14) *	72.46 (\pm 2.21)	71.04 (\pm 3.50)	71.45 (\pm 1.91)
CosmosQA	71.79 (\pm 1.55)	71.45 (\pm 3.02)	72.60 (\pm 1.85)	73.12 (\pm 2.15)
SQuAD	72.02 (\pm 1.88)	72.60 (\pm 2.07)	71.47 (\pm 3.24)	72.08 (\pm 2.25)
QQP	72.49 (\pm 2.79)	71.85 (\pm 2.98)	76.10 (\pm 1.26) *	71.30 (\pm 2.19)
N/A (Baseline)	71.07 (\pm 2.60)			

any target task’s performance on CR cases.

We show a pronoun error analysis in Appendix E to better understand the ALSC-CR improvements.

3.3 Evaluating Coreference Ability

Performing well on ALSC-CR requires strong CR ability, as CR associates the aspect with its sentiment. To verify that the improvement in Section 3.2 is attributable to the ALSC model’s improved CR ability, we estimate the CR ability by evaluating on DPR. Since we have an ALSC model for each random seed used for training (Section 3.2), we run DPR evaluation on the ALSC random seed model with the highest ALSC-CR val set performance.

DPR involves predicting the antecedent of a given pronoun. This is precisely the ability required for good performance on ALSC-CR (which contains only definite pronoun cases), making DPR ideal to measure the CR ability of models. Other CR datasets like OntoNotes (Hovy et al., 2006), Winograd Schema Challenge (WSC) (Levesque et al., 2012) and WinoGrande (Sakaguchi et al., 2021) are not as suitable as DPR because DPR only focuses on definite pronouns, which is the ability we are interested in. Similarly, DPR is also the only CR dataset suitable for auxiliary training, but the size makes this infeasible as discussed in Section 3.2.

We use a DPR variant for generative models where input is of the form: "Humans were afraid of robots as *they* were strong.", and the objective is to predict what the highlighted pronoun (*they*) is referring to (Raffel et al., 2019).

Evaluating ALSC models on DPR (Table 5) confirms that the ALSC-CR performance gains may be attributable to the improved CR ability of the model due to auxiliary fine-tuning. Experiments show that Commongen and QQP fine-tuned models show a drastically improved (and statistically significant) CR ability of up to \sim 16%. This explains their improved ALSC-CR performance. Using CosmosQA, we see a statistically significant \sim 5% deterioration

in CR ability which does not lead to statistically significant changes in ALSC-CR performance.

Table 5: CR ability of top performing models (Sec 3.2) measured using DPR. Statistically significant improvement(*) and deterioration(\dagger) from baseline marked. Best bolded, 2^{nd} best underlined.

Aux Task	Aux Frac.	Mean F1 (\pm Std. Dev)
N/A (Baseline)	0	59.28 (\pm 8.82)
Commongen	0.1	<u>75.77</u> (\pm 1.68)*
CosmosQA	1.0	54.55 (\pm 7.19) \dagger
SQuAD	0.2	62.91 (\pm 6.77)
QQP	0.5	76.36 (\pm 2.16)*

4 Related Work

Prior work notes CR to be important to ABSA and similar tasks (Kobayashi and Malon, 2022; Atkinson and Escudero, 2022). Ding and Liu (2010) use aspect sentiments for performing CR, demonstrating a correlation between CR and sentiment classification. De Clercq and Hoste (2020); De Bruyne et al. (2022) examine CR for detecting aspects from related reviews or images, for the reviews lacking explicit aspects. Instead, we consider an LLM’s intra-sentence CR ability, considering only reviews with explicit aspects as having an aspect is critical to ALSC. Mai and Zhang (2020) use CR in aspect extraction, but only for identifying duplicate references among proposed aspects. Varghese and Jayasree (2013) use CR to solve their dependency parser’s inability to correctly associate opinion words with pronouns. In our work, we consider the CR problem in end-to-end state-of-the-art ALSC LLM models. Chen et al. (2020) improve BERT LLM’s CR ability for opinion-mining, using a method relying on external knowledge bases.

5 Conclusion

Since real-world reviews vary widely, we need ALSC models which can handle various kinds of reviews, including those requiring CR. Although LLMs generally perform well on ALSC, our ex-

periments provide evidence that LLMs can have poor performance on ALSC reviews requiring CR ability. We show that this problem can be alleviated by fine-tuning with certain auxiliary tasks before fine-tuning on the target tasks. Our framework for evaluating and improving an LLM’s performance on CR cases can be applied for other tasks as well. Such a framework is critical for developing any model deployed in the real world. In the future, we will explore if auxiliary training can reduce the target task training that is needed for CR cases.

Limitations

- Even though we have successfully demonstrated a framework to handle CR-containing reviews by using auxiliary fine-tuning, we have not found which auxiliary tasks to definitively use for target tasks other than ALSC. The auxiliary task must be found using the framework proposed in our work.

- Our test set is composed of ~350 manually identified examples are guaranteed to require CR ability. However, it is common for ALSC datasets to be small. The bench-marking datasets Twitter, Lap14, Rest16 and Rest15 all have ~500-600 aspects for analysis (Zhang et al., 2019) which is close to our dataset. To reduce the variability due to a relatively small test set, we use multiple random seeds for robustness (Clark et al., 2020).

Due to the specific problem we are targeting, it is difficult to create more examples than this using existing sources. During qualitative analysis, we had considered many ALSC datasets (SemEval datasets, Twitter, MAMS) but found that the CR problem was most pronounced in the restaurant domain (Rest16, MAMS). Example: laptop reviews rarely use explicit aspects (Pontiki et al., 2014), leading to few CR cases in Lap14 dataset.

- Ours is the first work to demonstrate this CR problem in language models, thus there are few benchmarks against which we can compare our solution.
- We use the T5-large LLM for our experiments which requires a significant amount of computational resources for training. This leads to a high cost both financially and environmentally (Strubell et al., 2019).

References

- John Atkinson and Alex Escudero. 2022. [Evolutionary natural-language coreference resolution for sentiment analysis](#). *International Journal of Information Management Data Insights*, 2(2):100115.
- Jiahua Chen, Shuai Wang, Sahisnu Mazumder, and Bing Liu. 2020. [A knowledge-driven approach to classifying object and attribute coreferences in opinion mining](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1616–1626, Online. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). *arXiv preprint arXiv:2003.10555*.
- Luna De Bruyne, Akbar Karimi, Orphee De Clercq, Andrea Prati, and Veronique Hoste. 2022. [Aspect-based emotion analysis and multimodal coreference: A case study of customer comments on adidas Instagram posts](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 574–580, Marseille, France. European Language Resources Association.
- Orphee De Clercq and Veronique Hoste. 2020. [It’s absolutely divine! can fine-grained sentiment analysis benefit from coreference resolution?](#) In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 11–21, Barcelona, Spain (online). Association for Computational Linguistics.
- Xiaowen Ding and Bing Liu. 2010. [Resolving object and attribute coreference in opinion mining](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 268–276, Beijing, China. Coling 2010 Organizing Committee.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. [OntoNotes: The 90% solution](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Cosmos QA: Machine reading comprehension with contextual commonsense reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Qingnan Jiang, Lei Chen, Ruifeng Xu, Xiang Ao, and Min Yang. 2019. [A challenge dataset and effective models for aspect-based sentiment analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the*

- 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 6280–6285, Hong Kong, China. Association for Computational Linguistics.
- Hideo Kobayashi and Christopher Malon. 2022. [Analyzing coreference and bridging in product reviews](#). In *Proceedings of the Fifth Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 22–30, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020. [CommonGen: A constrained text generation challenge for generative commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1823–1840, Online. Association for Computational Linguistics.
- Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*. Morgan & Claypool Publishers.
- Quan Liu, Hui Jiang, Zhen-Hua Ling, Xiaodan Zhu, Si Wei, and Yu Hu. 2017. Combing context and commonsense knowledge through neural networks for solving winograd schema problems. In *2017 AAAI Spring Symposium Series*.
- Nicholas Lourie, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Unicorn on rainbow: A universal commonsense reasoning model on a new multitask benchmark. *AAAI*.
- Deon Mai and Wei Emma Zhang. 2020. [Aspect extraction using coreference resolution and unsupervised filtering](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 124–129, Suzhou, China. Association for Computational Linguistics.
- Jason Phang, Thibault Févry, and Samuel R Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- 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, 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.
- Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. [Intermediate-task transfer learning with pretrained language models: When and why does it work?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5231–5247, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Altaf Rahman and Vincent Ng. 2012. [Resolving complex cases of definite pronouns: The Winograd schema challenge](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 777–789, Jeju Island, Korea. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and policy considerations for deep learning in NLP](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy. Association for Computational Linguistics.
- Duyu Tang, Bing Qin, and Ting Liu. 2016. [Aspect level sentiment classification with deep memory network](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 214–224, Austin, Texas. Association for Computational Linguistics.
- Yuanhe Tian, Guimin Chen, and Yan Song. 2021. [Aspect-based sentiment analysis with type-aware graph convolutional networks and layer ensemble](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2910–2922, Online. Association for Computational Linguistics.

Raisa Varghese and M Jayasree. 2013. Aspect based sentiment analysis using support vector machine classifier. In *2013 international conference on advances in computing, communications and informatics (ICACCI)*, pages 1581–1586. IEEE.

Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2019. [Can you tell me how to get past sesame street? sentence-level pretraining beyond language modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4465–4476, Florence, Italy. Association for Computational Linguistics.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Hang Yan, Junqi Dai, Tuo Ji, Xipeng Qiu, and Zheng Zhang. 2021. [A unified generative framework for aspect-based sentiment analysis](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2416–2429, Online. Association for Computational Linguistics.

Chen Zhang, Qiuchi Li, and Dawei Song. 2019. [Aspect-based sentiment classification with aspect-specific graph convolutional networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4568–4578, Hong Kong, China. Association for Computational Linguistics.

Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2021. [Towards generative aspect-based sentiment analysis](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 504–510, Online. Association for Computational Linguistics.

Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2022. A survey on aspect-based sentiment analysis: Tasks, methods, and challenges. *arXiv preprint arXiv:2203.01054*.

A Hyperparameters

Learning rates for both auxiliary fine-tuning and ALSC training steps are picked from $\{5e-4, 1e-4, 5e-5\}$ and $\{1e-3, 5e-4, 1e-4\}$ respectively, after running for three random seeds and selecting the rates giving max F1 score for their respective validation dataset. For auxiliary fine-tuning, the learning rates for all auxiliary tasks were found to be $1e-4$, except for SQuAD with Aux Fraction as 1.0 for which we found learning rate as $5e-5$. For ALSC target task training, the learning rate was found to be $5e-4$ in all cases except when using CommonGen task for fine tuning with Aux Fraction as 0.1 for which we found learning rate as $1e-4$.

Batch size for training is taken as 16 to maximize GPU utilization. We train for 30 epochs to allow for convergence, while using an early stopping mechanism.

B Model Details

For our LLM, we use the T5-large implementation on Huggingface.³

C Dataset Cleanup

Following existing work (Tang et al., 2016; Tian et al., 2021) we disregard reviews with no aspects, and also the aspects labeled as having "conflict" sentiment polarity to prevent a class imbalance problem due to low count of "conflict" class.

D Dataset Details

Here we present some more details of the ALSC-CR dataset. The aspect polarity distribution is presented in Table 6. Note that it is possible to have multiple pronouns in each of the CR cases.

The sentiment distribution of ALSC-CR test set is shown in Table 7.

For constructing ALSC-CR, we use standard ALSC datasets (MAMS and Rest16). MAMS’s original train set along with data from Rest16 train set is used for training. For validation, we use the original validation sets from MAMS and Rest16, in addition to Pronoun cases from MAMS test and Rest16. The composition of the validation dataset is such that we use minimal Pronoun cases for validation while having sufficient CR cases for testing. Details of the composition of ALSC-CR are shown in Table 9.

³<https://huggingface.co/t5-large>

Table 6: Sentiment polarity distribution in ALSC-CR dataset

Partition	Polarity		
	Positive	Negative	Neutral
Train	4,279	3,065	5,090
Validation	337	222	330
Test	178	122	46

Table 7: Pronoun distribution in ALSC-CR test set, which has only CR cases

Pronoun	Count
it	132
which	59
they	54
he	24
who	19
she	17
their	14
them	12
its	10
his	10
there	10
him	5
her	5
hers	0

E Error Analysis by Pronoun

We analyze the errors and improvements seen for individual pronouns (in reviews) when ALSC-CR is evaluated with different ALSC models. Since a few pronouns have very low counts as per Table 7, we only analyze the ones which have count greater than 15.

For all pronouns analyzed, we find improvements in prediction accuracy for the models fine-tuned with auxiliary tasks, compared to the baseline model which has no auxiliary fine-tuning. Results are shown in Table 8.

F Training Prompts

We present the training prompts used in Table 10.

G Visualising Auxiliary Training Results

In Figure 1, we visually show the performance of auxiliary trained models on ALSC-CR (same results as Table 4). We can see that there is little correlation between the auxiliary dataset fraction and the mean F1 performance, making it necessary to explore various fraction settings.


H Training Details

For fine tuning the T5-large model, we use 1 NVIDIA V100 GPU, 6 CPU cores with 4 GB mem-

Table 8: Error Analysis of ALSC models by pronoun distribution. Model Accuracy% presented by Pronoun. Highest scores bolded. 2nd highest underlined. Pronouns with count less than 15 (as per Table 7) are not analyzed.

Pronoun	Baseline	Commongen 0.1	QQP 0.5
it	65.91	<u>68.18</u>	71.21
which	74.58	83.05	<u>77.97</u>
they	72.22	79.63	<u>77.78</u>
he	70.83	75.0	<u>70.83</u>
who	84.21	94.74	94.74
she	<u>88.24</u>	94.12	<u>88.24</u>
their	64.29	<u>78.57</u>	<u>78.57</u>
them	75.0	75.0	75.0
its	80.0	70.0	90.0
his	100.0	100.0	100.0
there	60.0	70.0	60.0
him	60.0	60.0	60.0
her	100.0	100.0	80.0
hers	N/A	N/A	N/A

Figure 1: Performance of ALSC models with aux training on ALSC-CR dataset.



ory per core. We run training jobs with a 71 hour time limit.

Table 9: Detailed ALSC-CR dataset composition.

Partition	Size	Composition
Train	12,434	MAMS Train (#count = 11,186) + Rest16 Train (Non Pronoun) (#count = 1,248)
Val	889	15% of (MAMS Test (Pronoun) + Rest16 Train/Val/Test (Pronoun)) + 50% of (MAMS Val + Rest Val (Non Pronoun)) [Here, MAMS #count = 746, Rest16 #count = 143]
Test	346	MAMS Test (CR) (#count = 124) + Rest16 Train/Val/Test (CR cases) (#count = 222)

Table 10: Details of T5 training prompts used for auxiliary and target tasks.

Task	Training Prompt
ALSC-CR	get sentiment: [sentence, aspect]
ALSC-Regular	get sentiment: [sentence, aspect]
DPR	Get antecedent: [sentence]
CommonGen	generate a sentence with: [concepts]
CosmosQA	question: [question] answer_0: [ans_0] answer_1: [ans_1] answer_2: [ans_2] answer_3: [ans_3] context: [context]
SQuAD	question: [question] context: [context]
QQP	qqp question1: [question_1] question2: [question_2]