

nicolay-r at SemEval-2024 Task 3: Using Flan-T5 for Reasoning Emotion Cause in Conversations with Chain-of-Thought on Emotion States

Nicolay Rusnachenko
Newcastle Upon Tyne, England
rusnicolay@gmail.com

Huizhi Liang
Newcastle University
Newcastle Upon Tyne, England
huizhi.liang@ncl.ac.uk

Abstract

Emotion expression is one of the essential traits of conversations. It may be self-related or caused by another speaker. The variety of reasons may serve as a source of the further emotion causes: conversation history, speaker’s emotional state, etc. Inspired by the most recent advances in Chain-of-Thought, in this work, we exploit the existing three-hop reasoning approach (THOR) to perform large language model instruction-tuning for answering: emotion states (THOR_{STATE}), and emotion caused by one speaker to the other (THOR_{CAUSE}). We equip THOR_{CAUSE} with the reasoning revision (RR) for devising a reasoning path in fine-tuning. In particular, we rely on the annotated speaker emotion states to revise reasoning path. Our final submission, based on Flan-T5_{base} (250M) and the rule-based span correction technique, preliminary tuned with THOR_{STATE} and fine-tuned with THOR_{CAUSE-RR} on competition training data, results in 3rd and 4th places ($F1_{\text{proportional}}$) and 5th place ($F1_{\text{strict}}$) among 15 participating teams. Our THOR implementation fork is publicly available: <https://github.com/nicolay-r/THOR-ECAC>

1 Task Overview

Extracting potential causes that lead to emotion expressions in text is the crucial aim of Emotion Cause Extraction (ECE) domain (Xia and Ding, 2019). In particular, the SemEval-2024 Task 3 (Wang et al., 2024) is aimed at emotion-cause pair analysis in conversations from the sitcom *Friends*. The conversations are organized into Emotion-Cause-in-Friends dataset (Wang et al., 2023) and includes the JSON-formatted training (TRAIN_{json}) and evaluation (TEST_{json}) parts. The authors propose 6 emotion classes to annotate: (i) speaker emotion states, and (ii) emotion caused by one utterance to the other. These classes are: $E = \{\text{SURPRISE, SADNESS, JOY, DISGUST, FEAR, ANGER}\}$,

and NEUTRAL for absence of emotion. We denote $E' = E \cup \{\text{NEUTRAL}\}$ as a complete set.

Among the several subtasks of ECAC-2024, in this paper we focused on *Subtask 1*: textual emotion-cause pair extraction in conversations. In this subtask, each conversation represents a list of utterances. Every utterance (u) yields the following: utterance text (u_{text}), speaker name (u_{speaker}), emotion state ($u_{\text{state}} \in E'$), and ID (u_{id}). The annotation of the emotion cause pairs represents a list $P = [p_1 \dots p_{|P|}]$, in which each pair $p \in P$ is a labeled source-target¹ tuple $p = \langle u^{\text{src}}, u^{\text{tgt}}, e_c \rangle$, where $e_c \in E$.

We initiate our studies by analyzing the training data (TRAIN_{json}) for the subject of annotated emotion-cause pairs $\langle u^{\text{src}}, u^{\text{tgt}} \rangle$ in it, and report:

1. Quantitative statistics of the mentioned emotion-cause pairs (Table 1);
2. Distance statistics (in utterances) between u^{src} and u^{tgt} (Table 2);
3. Distribution statistics between speaker state (u_{state}) and emotion *speaker causes* ($e^{u \rightarrow *}$) (Table 3).

According to the Table 2, most emotion was found to be caused by such utterances u^{src} that are the same as or mentioned before u^{tgt} ($\delta \geq 0$). Therefore, given $\langle u^{\text{src}}, u^{\text{tgt}} \rangle$ we denote its context $X = \{u^1 \dots u^k\}$ as a *history* of the past $k - 1$ utterances of u^{tgt} , where $u^{\text{tgt}} = u^k \in X$, $u^{\text{src}} \in X$. **Task definition:** Given an emotion-causing utterance pair within context $\langle u^{\text{src}}, u^{\text{tgt}}, X \rangle$ answer the emotion $e_c \in E'$ caused by u^{src} towards u^{tgt} .

2 Methodology

We propose a two-stage training mechanism for performing instruction-tuning on large language models (LLMs), aimed at accurately inferring of

¹Spans-prediction is beyond the scope of our methodology.

Parameter	Value
Conversations (total)	1374
Emotion causes pairs per conversation	6.46
Emotion causes pairs in annotation (total)	8879
Self-cause per conversation (% from total)	51.86%
Self-cause by different utterance (% from total)	12.83%

Table 1: Quantitative statistics of the emotion-cause pairs in the competition training data (TRAIN_{json})

Parameter	future	past				
		0	1	2	3	4
$\delta = u_{id}^{tgt} - u_{id}^{src}$	< 0					
Causes count	377	4605	2759	810	332	160
Average per δ	0.12	3.35	2.01	0.59	0.24	0.12
Covering (%)	-	51.9	82.9	92.1	95.8	97.6

Table 2: Distance statistics (δ) (in utterances) between source (u^{src}) and target (u^{tgt}) of emotion-cause pairs in the competition training data (TRAIN_{json})

$u_{state} \setminus e^{u \rightarrow *}$	JOY	SUR	ANG	SAD	DIS	FEA
total	2653	2092	1984	1336	518	296
JOY	.89	.06	.03	.01	.01	.00
SURPRISE	.07	.78	.07	.03	.03	.02
ANGER	.01	.07	.83	.06	.02	.02
SADNESS	.02	.09	.06	.81	.01	.01
DISGUST	.03	.07	.14	.06	.70	.01
FEAR	.02	.13	.08	.05	.04	.68
NEUTRAL	.24	.38	.22	.08	.04	.03

Table 3: . Distribution statistics between speaker state (u_{state}) and emotion *speaker causes* ($e^{u \rightarrow *}$) in the competition training data (TRAIN_{json}); values in each row are normalized

the task answers. Given triplet $\langle u^{src}, u^{tgt}, X \rangle$ of emotion-cause pair $\langle u^{src}, u^{tgt} \rangle$ in context X , the proposed mechanism aims at LLM instruction-tuning, in order to answer $e \in E'$ that refers to:

STAGE 1: emotion state u_{state}^{tgt} ;

STAGE 2: emotion cause by u^{src} to u^{tgt} .

Therefore, for emotion-cause pairs extraction we use the STAGE 2 towards the model tuned in STAGE 1 to infer $e_c \in E'$ caused by u^{src} towards u^{tgt} .

Instead of directly asking LLM the final result at each stage, we exploit the Chain-of-Thought (CoT) concept in the form of the Three-hop Reasoning (THOR) framework (Hao et al., 2023). We believe that LLM can infer the span that conveys emotion and opinion about it before answering $e \in E'$. Figure 1 illustrates the proposed training methodology, empowered by the CoT prompting. We refer to the instruction-tuning mechanisms of the STAGE 1 and STAGE 2 as THOR_{STATE} and THOR_{CAUSE} respectively.

2.1 Chain-of-Thought Prompting

We adopt the THOR framework (Hao et al., 2023) in LLM fine-tuning with the prompt templates adapted for emotion-cause pair analysis in conversations. We define the intermediate *span* (s) and latent *opinion* expression (o). With $C_i, i \in \overline{1..3}$ we denote the prompts that wrap the content in the input context. The construction of stages is as follows.

THOR_{STATE} This is a STAGE 1 of the proposed training methodology, aimed at preliminary LLM instruction-tuning. Given $\langle u^{tgt}, X \rangle$, we apply the following three steps to infer $u_{state}^{tgt} = e'_1 \in E'$:

Step 1: $s'_1 = [C_1(X)$, which text spans are possibly causes emotion on u_{text}^{tgt} ?]

Step 2: $o'_1 = [C_2(C_1, s'_1)$. Based on the common sense, what is the implicit opinion towards the mentioned text spans that causes emotion on u_{text}^{tgt} , and why?]

Step 3: $e'_1 = [C_3(C_2, o'_1)$. Based on such opinion, what is the emotion state of u_{text}^{tgt} ?]

where s'_1 could be interpret as $s'_1 = \text{argmax } p(s_1|X, u_{text}^{tgt})$, latent opinion o'_1 as $o'_1 = \text{argmax } p(o_1|X, u_{text}^{tgt}, s'_1)$, and the final answer e'_1 noted as: $e'_1 = \text{argmax } p(e_1|X, u_{text}^{tgt}, s'_1, o'_1)$.

THOR_{CAUSE} This is a STAGE 2 of the proposed methodology, based on emotions-cause pairs. We use this stage for (i) *fine-tuning* and (ii) task result *inferring* purposes. Given context $\langle u^{src}, u^{tgt}, X \rangle$ we omit² $u^{tgt} \in X$ from the input parameters by referring to it as «end of the conversation». We apply the following steps to infer $e'_2 \in E'$ caused by u^{src} to u^{tgt} :

Step 1: $s'_2 = [C_1(X)$, which specific text span of u_{text}^{src} is possibly causes emotion?]

Step 2: $o'_2 = [C_2(C_1, s'_2)$. Based on the common sense, what is the implicit opinion towards the cause of mentioned text span of u_{text}^{src} , and why?]

Step 3: $e'_2 = [C_3(C_2, o'_2)$. Based on such opinion, what is the emotion caused by source towards the last conversation utterance?]

²To reduce the problem statement to the one for which THOR was originally designed (Pontiki et al., 2016)

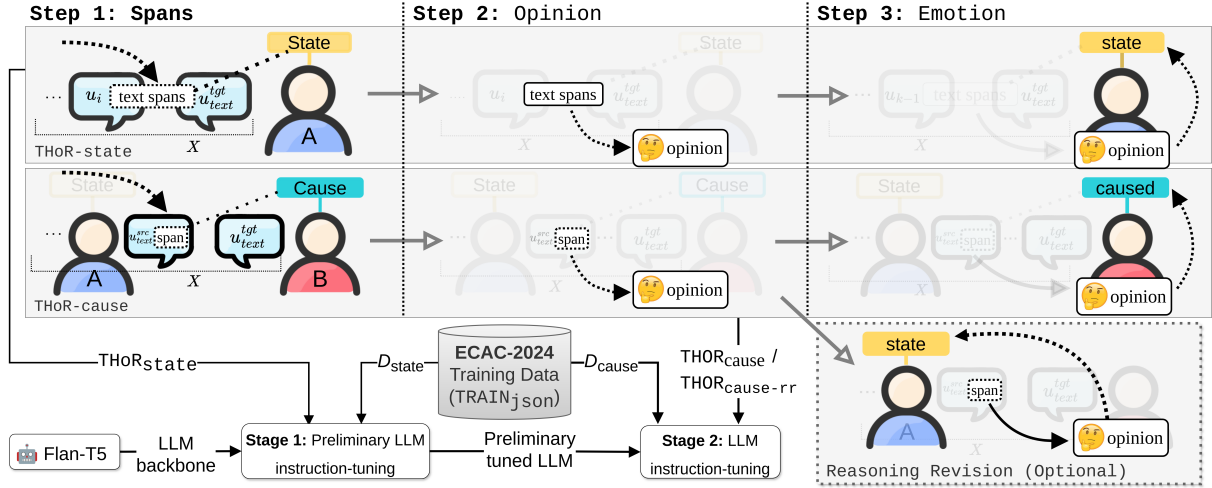


Figure 1: Two-stage LLM tuning methodology for inferring emotion caused by u^{src} towards u^{tgt} in context X by adapting THOR (Hao et al., 2023) to reason and answer: (i) u^{tgt} (THOR_{STATE}), and (ii) emotion caused by u^{src} towards u^{tgt} (THOR_{CAUSE}), optionally enhanced by Reasoning-Revision and by predicting u^{state} (THOR_{CAUSE-RR}).

where s'_2 could be interpreted as $s'_2 = \text{argmax } p(s'_2|X, u^{src}_{text})$, opinion o' could be interpreted as $o'_2 = \text{argmax } p(o_2|X, u^{src}_{text}, s'_2)$, and the final answer e'_2 noted as: $e'_2 = \text{argmax } p(e_2|X, u^{src}_{text}, s'_2, o'_2)$.

2.2 Reasoning Revision with Supervision

During the LLM instruction-tuning process with the THOR, it is possible to devise a reasoning path. Technically, at each step of the chain we have all the necessary information to query our model with the final answer. With the following approach, we believe in a better model alignment on state-cause dependency (Table 3): speakers are likely to cause an emotion, similar to their states³. To revise this knowledge, in this paper, we impute the following prompt to support our opinion O , obtained at the end of the THOR_{CAUSE} step 2 (Fig. 1):

Step 3.1: $u^{state} = [C_3(C_2, o'_2)]$, Based on such opinion, what is the emotion state of u^{src}_{text} ?

Due to the definition of the task, we believe in the correctness of this knowledge within the emotion cause task. Once step 3.1 is embedded, the result answer $e'_2 \in E'$ in THOR_{CAUSE} from the step 3 could be reinterpreted as $e'_2 = \text{argmax } p(e_2|X, u^{src}_{text}, s'_2, o'_2, u^{state})$. We refer to this setup as THOR_{CAUSE-RR}.

³Except NEUTRAL speaker state (Table 3)

3 Datasets and Experiential Setup

We adopt textual resources provided by the competition organizers (Wang et al., 2024): training (TRAIN_{json}) and evaluation (TEST_{json}) data. Within TRAIN_{json}, for each conversation, we rely on (i) speakers *emotion states*, and (ii) *emotion causes* annotation to compose the datasets D_{state} and D_{cause} , respectively. Each dataset represent a list of tuples $t = (u, X, L)$, where u is an utterance of the conversation context $X = \{u^1 \dots u^k\}$, and L is a list of emotion labels, defined as:

- $L = [u^{state}_k]$ in the case of D_{state} ($u^{state}_k \in E'$)
- $L = [u^{state}, e^u]$ in the case of D_{cause} , where e^u is emotion expressed by u towards u^k , or NEUTRAL otherwise ($e^u \in E'$)

D_{state} represent entries of all possible utterances in all conversations with their emotional states $u_{state} \in E$. For the particular utterance u , we consider its context as $X_u = \{u' : u_{id} - u'_{id} \leq k\}$.

D_{cause} includes all possible pairs $\langle u^{src}, u^{tgt} \rangle$, where $u^{src}_{id} \leq u^{tgt}_{id}$, and $u^{tgt}_{id} - u^{src}_{id} \leq k$. For the particular pair, we compose the related context (X') as follows: $X' = \{u' : u^{tgt}_{id} - u'_{id} \leq k\}$. For each pair, we assign $e \in E$ if the pair is present in conversation annotation and NEUTRAL otherwise. We rely on the analysis in Table 2 to limit the number of pairs, as well as the size of the context. We set $k = 3$ to cover 95.8% emotion-cause pairs. We also cover the case of emotions caused from within the same utterance (59.5%, see Table 1). As for emotions caused by the same speaker of

Source	TRAIN _{json}		TEST _{json}
Part	train	dev	test
D_{state} (total)	12144	1475	
NEUTRAL	5299	630	.
JOY	2047	254	.
SURPRISE	1656	184	.
ANGER	1423	192	.
SADNESS	1011	136	.
DISGUST	372	42	.
FEAR	336	37	.
D_{cause} (total)	30445	3612	15794
NEUTRAL	23750	2765	15794
JOY	2111	279	–
SURPRISE	1725	202	–
ANGER	1307	174	–
SADNESS	932	120	–
DISGUST	387	47	–
FEAR	233	25	–

Table 4: Statistics of the composed datasets D_{state} and D_{cause} from the publicly available competition data, for the two training methodology stages respectively; statistics is listed for $k = 3$.

other utterance, we assess that excluding this type of pairs (12.83%, according to Table 1), results in $\approx 23\%$ pairs reduction of D_{cause} and hence reduces training time. Therefore, the result D_{cause} excludes pairs of this type in `train`, `dev` and `test` parts.

Table 4 lists the statistics of the composed resources. We use the 9:1 proportion for TRAIN_{json} to compose `train` and `dev`, respectively. To represent $X \in t$, we concatenate its representation of utterances. For each utterance $u \in X$, we use the following formatting template: « $u_{speaker} : u_{text}$ ». To represent utterance $u \in t$, we refer to u_{text} . For each $l \in L$ formatting, we utilize its lowercase text value. The implementation details for the datasets preparation are publicly available.⁴

Setup. We follow the publicly available framework setups (Hao et al., 2023) and adopt encoder-decoder style instructive Flan-T5⁵ as our backbone LLM for the proposed methodology. We experiment with a 250M (base) version. For evaluations on `dev`, we adopt the F1-measure for E' , denoted as $F1(E')$. The evaluation on `test` assessed with the set of $F1$ -metrics, provided by the competition organizers (details in Section 4). We consider the instruction-tuning of the Flan-T5 model with the following techniques: conventional PROMPT, THOR (Section 2.1), and THOR_{CAUSE} with reasoning revision (Section 2.2). To conduct the experiment, we rent a server with a single NVIDIA A100 GPU (40GB). We set temperature 1.0, learning rate

⁴<https://github.com/nicolay-r/SemEval2024-Task3>

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

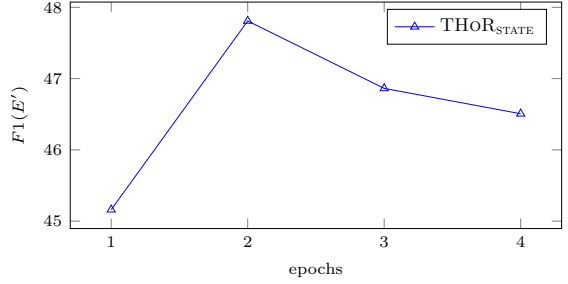


Figure 2: Result analysis of the preliminary fine-tuning of Flan-T5_{base} on D_{state} dev using THOR_{STATE} technique per epoch by $F1(E')$

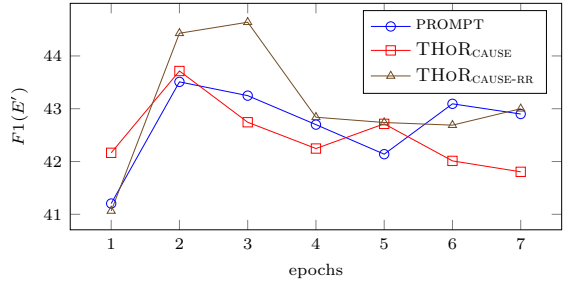


Figure 3: Flan-T5_{base}[†] fine-tuning results comparison by $F1(E')$ on D_{cause} dev part per each epoch across fine-tuning techniques: PROMPT, THOR_{CAUSE}, and THOR_{CAUSE-RR}.

$2 \cdot 10^{-4}$, optimizer AdamW (Loshchilov and Hutter, 2017), BATCH-SIZE of 32.

For the PROMPT technique, we use the template « $C_1(X). I(u). Choose from E'$ », where $I(u)$ corresponds to the instruction. For D_{cause} we use $I(u) = \llcorner What emotion causes u_{text} towards the last conversation utterance? \rceil$

4 Experiments

Stage 1. Figure 2 illustrates the analysis of the $F1$ on `dev` part during the preliminary tuning of Flan-T5_{base} on D_{state} .⁶ We investigate the overfitting after 2 epochs of training. The best state, obtained at the end of the epoch #2 with the $F1(E') = 47.81$ on the D_{state} -`dev` part, has been selected. In further, we refer to this model as Flan-T5_{base}[†].

Stage 2 Figure 3 provides a comparative analysis of different fine-tuning techniques. As at the pre-training stage, we investigate the ability to learn task emotion states 2-3 training epoch, followed by overfitting. Switching from PROMPT to THOR_{CAUSE-RR} technique, we investigate the improvement by 2.5% percent by $F1(E')$ on the `dev`

⁶We left the comparison with other pre-training techniques listed in 3 out of scope of this paper due to alignment with the CoT concept in STAGE 2.

Source	dev		test		
Model	$F1(E')$	$F1_s^w$	$F1_p^w$	$F1_s$	$F1_p$
PROMPT					
FT5 _{base} [‡]	43.51	9.68	22.27	10.05	22.21
THOR_{CAUSE}					
FT5 _{base} [‡]	43.72	–	–	–	–
THOR_{CAUSE-RR}					
FT5 _{base} [‡]	44.64	9.74	23.54	10.33	23.94
THOR_{CAUSE-RR} + Algorithm-based Spans Correction					
FT5 _{base} [‡]	44.64	12.86	24.28	13.26	24.13

Table 5: Evaluation results for Flan-T5_{base}[‡] on dev and test parts of the D_{cause} dataset; the results of the final submission are highlighted in gray

part of the D_{state} dataset. We refer to the best fine-tuned versions as Flan-T5_{base}[‡], separately per each fine-tuning technique in Table 5 (dev column).

The official evaluation includes the following $F1$ measures: (i) weighted averaged $F1_*^w$ / non-weighted ($F1_*$), and (ii) strict ($F1_s$) / not-strict ($F1_p$) towards predicted spans. To form the submissions for official evaluation, the following span corrections approaches were used: (i) punctuation terms⁷ exclusion from utterance prefixes and suffixes (by default), and (ii) algorithm-based (Section 4.1). Table 5 (test columns) illustrate the available results of T5_{base}[‡] in official evaluation.

Final submission represents the results of Flan-T5_{base}[‡] (THOR_{CAUSE-RR} technique), and application of algorithm-based spans correction.

4.1 Algorithm-based Spans Correction

Our methodology (Section 2) is limited on utterance level emotion cause prediction.⁸ We believe it is reflected in the relatively low results of $F1_s$ on the test dataset (see Table 5). Therefore, we analyze TRAIN_{json} and adopt a placeholder solution, aimed at enhancing the results by $F1_s$.

We apply a *rule-based approach* based on differences between the original utterance texts and their span annotations in the training data. Using TRAIN_{json}, we compose *prefix-* (V_p) and *suffix-* (V_s) vocabularies. For vocabulary entries, we select those that satisfy all of the following criteria: (i) the length of entry does not exceed 5 words, (ii) entry starts (in the case of V_s), or ends (in the case of V_p) with the punctuation sign⁷.

⁷We use `string.punctuation` preset in Python

⁸Technically it is possible to obtain spans (Section 2), however we could not investigate the practical valuty of the THOR_{CAUSE}-based Flan-T5_{base}[‡] responses from step #1.

Parameter	Value
Conversations (total)	2917
Emotion causes pairs in annotation	665
Average per conversation	4.39

Table 6: Quantitative statistics of the automatically extracted emotion-cause pairs by Flan-T5_{base}[‡] (THOR_{CAUSE-RR} technique) from the evaluation data (TEST_{json})

Parameter	past			
$\delta = u_{id}^{tgt} - u_{id}^{src}$	0	1	2	3
Causes count	1711	1012	148	46
Average per δ	2.57	1.52	0.22	0.07
Covering (%)	58.7	93.3	98.4	100.0

Table 7: Statistic of distances in utterances (δ) between source (u^{src}) and target (u^{tgt}) of emotion-cause pairs for automatically extracted emotion-cause pairs by Flan-T5_{base}[‡] (THOR_{CAUSE-RR} technique) from the evaluation data (TEST_{json})

$u_{state} \setminus e^{u \rightarrow *}$	JOY	SUR	ANG	SAD	DIS	FEA
JOY	.87	.08	.02	.01	.01	.00
SURPRISE	.09	.75	.06	.05	.03	.01
ANGER	.05	.14	.68	.08	.03	.01
SADNESS	.06	.11	.03	.76	.02	.02
DISGUST	.07	.11	.07	.05	.68	.01
FEAR	.00	.15	.09	.02	.00	.74
NEUTRAL	.36	.40	.07	.12	.03	.02

Table 8: Distribution statistics between speaker state (u_{state}) and emotion *speaker causes* ($e^{u \rightarrow *}$) for automatically extracted emotion-cause pairs by Flan-T5_{base}[‡] (THOR_{CAUSE-RR} technique) from the evaluation data (TEST_{json}); values in each row are normalized

$u_{state} \setminus e^{* \rightarrow u}$	JOY	SUR	ANG	SAD	DIS	FEA
JOY	.97	.01	.01	.00	.01	.00
SURPRISE	.04	.89	.04	.01	.01	.01
ANGER	.04	.05	.83	.05	.02	.01
SADNESS	.02	.02	.03	.89	.02	.01
DISGUST	.02	.04	.05	.07	.81	.01
FEAR	.00	.06	.07	.04	.03	.80
NEUTRAL	.60	.13	.03	.16	.05	.02

Table 9: Distribution statistics between speaker state (u_{state}) and emotion *caused on them* ($e^{* \rightarrow u}$), for automatically extracted emotion-cause pairs by Flan-T5_{base}[‡] (THOR_{CAUSE-RR} technique) from the evaluation data (TEST_{json}); values in each row are normalized

For each utterance text (u_{text}) that causes emotion, we compose an updated u'_{text} by applying: (1) correction of u_{text} prefixes with V_p , followed by (2) correction of suffixes from V_s for the results from (1). We alter u'_{text} in the case of $u'_{text} = \emptyset$. The algorithm 1 illustrates an implementation for the prefixes correction with V_p .⁹

⁹Implementation is publicly available in <https://github.com/nicolay-r/SemEval2024-Task3>

Algorithm 1 Emotion-cause prefixes correction for u_{text}

```
updated  $\leftarrow$  True
 $V_p' \leftarrow$  sorted  $V_p$  by decreased entry lengths in words
while  $u_{text} \neq \emptyset$  or updated do
  updated  $\leftarrow$  False
   $u_{text}' \leftarrow u_{text}$   $\triangleright$  Modified version of  $u_{text}$ 
  for  $v_p \in V_p'$  do
    if  $u_{text}'$  ends with  $v_p$  then
       $u_{text}' \leftarrow$  part of  $u_{text}'$  before  $v_p$ 
      updated  $\leftarrow$  True
      break
    end if
  end for
end while
```

4.2 Final Submission Analysis

We report the following emotion-cause pairs $\langle u^{src}, u^{tgt} \rangle$ analysis results for the Flan-T5_{base}‡ (THOR_{CAUSE-RR} technique, final submission):

1. Quantitative statistics of the extracted emotion-cause pairs (Table 6);
2. Distance statistics (in utterances) between u^{src} and u^{tgt} (Table 7);
3. Distribution statistics between speaker state (u_{state}) and the emotion *speaker causes* ($e^{u \rightarrow *}$) (Table 8);
4. Distribution statistics between speaker state (u_{state}) and emotion *caused on them* ($e^{* \rightarrow u}$) (Table 9).

According to the results in Table 8, we observe that the correlation between the state of the speaker u utterance (u_{state}) and the emotion it causes ($e^{u \rightarrow *}$) is **similar to** the related statistics on the competition training data (Table 3). We also investigate the alignment of the speaker states (u_{state}) with the emotion caused on them ($e^{* \rightarrow u}$) and the precision of the result varies between 80–97% (Table 9). The known source of misalignment is the case when emotion¹⁰ $e^{* \rightarrow u} \in E$ caused on u with $u_{state} = \text{NEUTRAL}$ (bottom row, Table 9).

5 Conclusion

In this paper, we present a Chain-of-Thought (CoT) methodology aimed at fine-tuning LLM for emotion state and cause extraction. We consider the problem of *emotion cause analysis in conversations* as a context-based problem with the mentioned utterance that causes emotion towards the last utterance in context. We devise our CoT for

¹⁰JOY especially, as the most frequently appearing class.

emotion causes and propose a reasoning revision methodology aimed at imputing the speaker emotion to support the decision on caused emotion. Our CoT represent a Three-hop Reasoning approach priority known as THOR. We apply this approach to fine-tune LLM and predict: (i) emotion state of the mentioned utterance, and (ii) emotion caused by mentioned utterance towards the last utterance in context. We experiment with the Flan-T5_{base} (250M) model fine-tuning using resources provided by task organizers. The application of CoT with reasoning revision allows us to improve the results by 2.5% (F1-measure) compared to prompt-based tuning. In further work, we expect to contribute with the: (i) analysis of larger models, and (ii) enhanced reasoning revision techniques, mentioned in the final submission analysis.

References

- Fei Hao, Li Bobo, Liu Qian, Bing Lidong, Li Fei, and Chua Tat-Seng. 2023. Reasoning implicit sentiment with chain-of-thought prompting. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1171–1182.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammed AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, et al. 2016. Semeval-2016 task 5: Aspect based sentiment analysis. In *ProWorkshop on Semantic Evaluation (SemEval-2016)*, pages 19–30. Association for Computational Linguistics.
- Fanfan Wang, Zixiang Ding, Rui Xia, Zhaoyu Li, and Jianfei Yu. 2023. Multimodal emotion-cause pair extraction in conversations. *IEEE Transactions on Affective Computing*, 14(3):1832–1844.
- Fanfan Wang, Heqing Ma, Rui Xia, Jianfei Yu, and Erik Cambria. 2024. **Semeval-2024 task 3: Multimodal emotion cause analysis in conversations**. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 2022–2033, Mexico City, Mexico. Association for Computational Linguistics.
- Rui Xia and Zixiang Ding. 2019. Emotion-cause pair extraction: A new task to emotion analysis in texts. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1003–1012.