arXiv:2406.02642v1 [cs.LG] 4 Jun 2024

E-ICL: ENHANCING FINE-GRAINED EMOTION RECOGNITION THROUGH THE LENS OF PROTOTYPE THEORY

Zhou Yang[1], **Zhaochun Ren**[2], **Chenglong Ye**[1], **Yufeng Wang**[1], **Haizhou Sun**[3], **Chao Chen**[4], **Xiaofei Zhu**[5], **Yunbing Wu**[1], **Xiangwen Liao**[1*]
[1]College of Computer and Data Science, Fuzhou University, Fuzhou, China;
[2]Leiden University, Leiden, The Netherlands; [3]H. Sun is with SmartMore;
[4] School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China;
[5]College of Computer Science and Technology, Chongqing University of Technology, Chongqing, China
{200310007, 211027083, 102102153, liaoxw}@fzu.edu.cn
z.ren@liacs.leidenuniv.nl zxf@cqut.edu.cn cha01nbox@gmail.com

ABSTRACT

In-context learning (ICL) achieves remarkable performance in various domains such as knowledge acquisition, commonsense reasoning, and semantic understanding. However, its performance significantly deteriorates for emotion detection tasks, especially fine-grained emotion recognition. The underlying reasons for this remain unclear. In this paper, we identify the reasons behind ICL's poor performance from the perspective of prototype theory and propose a method to address this issue. Specifically, we conduct extensive pilot experiments and find that ICL conforms to the prototype theory on fine-grained emotion recognition. Based on this theory, we uncover the following deficiencies in ICL: (1) It relies on prototypes (example-label pairs) that are semantically similar but emotionally inaccurate to predict emotions. (2) It is prone to interference from irrelevant categories, affecting the accuracy and robustness of the predictions. To address these issues, we propose an Emotion Context Learning method (E-ICL) on fine-grained emotion recognition. E-ICL relies on more emotionally accurate prototypes to predict categories by referring to emotionally similar examples with dynamic labels. Simultaneously, E-ICL employs an exclusionary emotion prediction strategy to avoid interference from irrelevant categories, thereby increasing its accuracy and robustness. Note that the entire process is accomplished with the assistance of a plug-and-play emotion auxiliary model, without additional training. Experiments on the fine-grained emotion datasets EDOS, Empathetic-Dialogues, EmpatheticIntent, and GoEmotions show that E-ICL achieves superior emotion prediction performance. Furthermore, even when the emotion auxiliary model used is lower than 10% of the LLMs, E-ICL can still boost the performance of LLMs by over 4% on multiple datasets.

1 INTRODUCTION

Achieving human-like intelligence necessitates endowing machines with the capability to perceive nuanced human emotions Poria et al. (2017); Picard (2000); Calvo & D'Mello (2010). Fine-grained emotion recognition (FER) Liew & Turtle (2016); Abdul-Mageed & Ungar (2017) aims to identify a variety of nuanced emotion categories, which is widely integrated into downstream tasks like empathetic dialogue Rashkin et al. (2019); Sabour et al. (2022); Li et al. (2022; 2020); Yang et al. (2023b); Zhao et al. (2022), sentiment analysis Wang et al. (2016); Schuff et al. (2017); Guzman & Maalej (2014), and emotional support systems Saha et al. (2021; 2022); Peng et al. (2022); Tu et al. (2022), yielding significant benefits.

* Corresponding author.

Early methods train small-scale models to identify fine-grained emotions in the query Kim et al. (2021); Majumder et al. (2020); Xie et al. (2019); Majumder et al. (2019); Ghosal et al. (2019). These methods are often confined to the emotions and knowledge within specific datasets, lacking flexibility and generalizability. Recently, in-context learning (ICL) has shown outstanding performance across a wide range of tasks by prompting large language models (LLMs) to comprehend the query and relevant demonstrations Rae et al. (2021); Liu et al. (2021); Yang et al. (2023a); Xiao et al. (2023); Liu et al. (2021); Rubin et al. (2021); Fu et al. (2022), exhibiting remarkable flexibility and generalizability. However, ICL struggles with fine-grained emotion recognition, suffering from poor performance Zhao et al. (2023); Schaaff et al. (2023); Yang et al. (2024); Qian et al. (2023). Concurrently, their performance declines further as the relevance of the demonstrations decreases Xu et al. (2024); Liu et al. (2021). For example, randomly selected demonstrations within the training dataset fare worse than semantically relevant ones, while demonstrations outside the training dataset perform even more poorly.

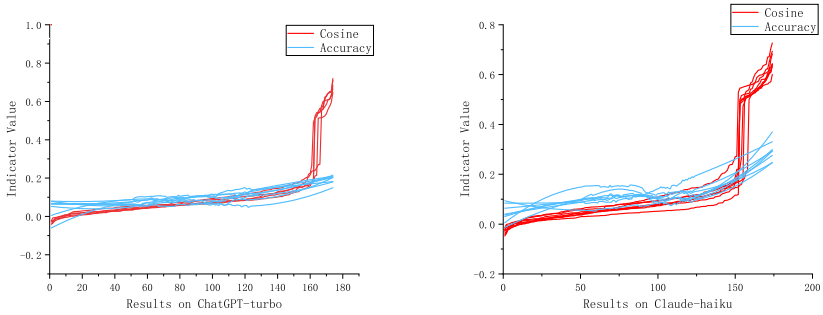

Figure 1: Results of the pilot experiment, where the red line represents the similarity between the example and the query, and the blue line denotes the ICL's emotion accuracy.

To investigate this phenomenon and its underlying reasons, we conduct pilot experiments on multiple fine-grained emotion datasets (The details are provided in Appendix A). Specifically, we construct multiple sets of examples and select queries with varying similarity levels for each set. We then employ the ICL method to predict the queries' emotion categories on different LLMs, i.e., ChatGPT-turbo and Claude-haiku. The experimental results are shown in Figures 1 (a) and (b), respectively. The results indicate that as the semantic similarity increases, the emotion accuracy of the query also continues to rise. Framing queries as samples and demonstrations as prototypes, we observe that ICL aligns with the core tenet of prototype theory Rosch (1978); Kamp & Partee (1995); Hampton (2006): the more similar the sample (query) is to the prototype (demonstrations), the more accurately the sample can be recognized.

From the perspective of prototype theory, ICL in FER can be reinterpreted as the following process: prompting LLMs to refer to the emotions of demonstrations, i.e., semantically similar example-label pairs, to predict the query's emotion categories. However, according to prototype theory, this process has two deficiencies: **(1) Referring to semantically similar example-label pairs are prone to causing emotion misunderstanding**. For instance, the example "I'm devastated about the news" is semantically similar to the query "I'm delighted about the news." Referring to the example, the query's emotion is likely to be misunderstood as negative. Meanwhile, the emotional labels of examples are annotated as static and unique, while human emotions are typically dynamic and diverse Larsen & McGraw (2011); Crivelli & Fridlund (2019); Trampe et al. (2015). For example, the example "That's so great" is labeled as "excited" in ICL, whereas a more complete emotion would be "joyful, excited." Referring to the former would ignore the emotion "joyful" in the example, while the latter would provide a more comprehensive emotional reference. (2) **ICL is susceptible to interference from irrelevant emotion categories during prediction.** FER involves numerous emotions that are difficult to distinguish, such as "joyful" and "excited". In predicting these emotions, ICL treats all emotions as candidate options and makes judgments without carefully considering these options. This hinders ICL's ability to discern the nuanced distinctions between emotions, resulting in inaccurate predictions. In contrast, prototype theory-based approaches prioritize more likely categories, such as those closer to the prototypes, when faced with similar scenarios Snell et al.

(2017); Ding et al. (2021); Nauta et al. (2021). By excluding irrelevant options, this approach allows for careful consideration of the more probable categories, resulting in more robust and accurate predictions.

In this paper, we propose Emotional In-Context Learning (E-ICL) for fine-grained emotion recognition. With the prototype theory perspective, E-ICL refers to emotionally similar examples with dynamic labels to make exclusionary emotion predictions. E-ICL first retrieves emotionally similar examples to avoid misunderstandings caused by semantically similar examples. It then assigns dynamic emotion labels to the constructed examples, comprehensively representing their emotional state. Finally, E-ICL refers to the prototypes, i.e., the constructed example-label pairs, and makes exclusionary emotion predictions to determine the query's emotion category. The exclusionary emotion predictions reduce the interference of implausible emotion categories during prediction, enabling more robust and accurate predictions. It is noteworthy that the entire process is accomplished with the assistance of a plug-and-play emotion-capable auxiliary model, without requiring any model training, further enhancing the flexibility and applicability of E-ICL.

We conduct extensive experiments on four fine-grained emotion recognition datasets: EDOS Welivita et al. (2021), Empathetic-Dialogues (ED) Rashkin et al. (2019), EmpatheticIntent (EI) Welivita & Pu (2020), and GoEmotions (GE) Demszky et al. (2020). The experimental results demonstrate that compared to ICL, E-ICL guides LLMs to perceive fine-grained emotions more accurately with the assistance of different emotion-capable auxiliary models. Furthermore, more analyses show that E-ICL exhibits stable performance across different auxiliary models and LLMs. Notably, even when the performance of the auxiliary model is 10% lower than that of LLMs, the proposed method still enhances LLMs with a 4% higher performance than ICL on multiple datasets, indicating its stable advantage.

In summary, our contributions are as follows: (1) We find that ICL aligns with prototype theory, and based on this theory, we uncover its limitations for fine-grained emotion recognition tasks. We then propose E-ICL to address this limitation. To the best of our knowledge, we are the first approach to explore ICL from this perspective. (2) We introduce a novel strategy for constructing example-label pairs by building emotionally similar examples with dynamic emotion labels, providing a new perspective on the demonstration construction for ICL. (3) We introduce an exclusionary emotion prediction strategy, inspired by the prediction strategy of prototype theory, to make more accurate and robust emotion predictions. (4) The experiments show that E-ICL exhibits a stable advantage in fine-grained emotion recognition across multiple datasets.

2 Related Work

Fine-grained Emotion Recognition. The goal of the fine-grained emotion recognition task (FER) is to detect subtle emotion categories in the query Liew & Turtle (2016); Abdul-Mageed & Ungar (2017). As emotions are primarily influenced by situational and cognitive factors Gross et al. (2014); Siemer et al. (2007); Moors et al. (2013), existing works have mainly explored these two aspects and can be divided into situation-based models and cognition-based models. Situation-based Models mainly detect the subtle emotions implied in the query, without considering additional cognition information. These Models have explored word-level emotions Li et al. (2020); Kim et al. (2021); Yang et al. (2023b); Wang et al. (2024), mixed emotions Majumder et al. (2020); Lin et al. (2019), and sentence-level emotions Xie et al. (2019); Majumder et al. (2019); Ghosal et al. (2019). Cognition-based models mainly enhance emotions through additional cognitive factors. These models have explored aspects such as emotion causes, commonsense knowledge, and personalized features. Both types of models have played an important role in the FER. However, these models are trained on specific datasets, limited by the corresponding data, and require certain computational resources and training time. In contrast, we explore the FER task through In-Context Learning, without consuming computational resources and training time.

In-Context Learning. In-Context Learning (ICL) improves LLMs' performance by learning from constructed demonstrations, circumventing the time and computational costs associated with fine-tuning. One part of ICL enhances LLMs by breaking down the reasoning steps of demonstrations into sub-steps and enabling LLMs to complete tasks by following these sub-steps Wei et al. (2022); Hendrycks et al. (2021); Kazemi et al. (2022). This type of ICL has demonstrated satisfactory results in tasks such as arithmetic, commonsense, and symbolic reasoning Rae et al. (2021). However, these

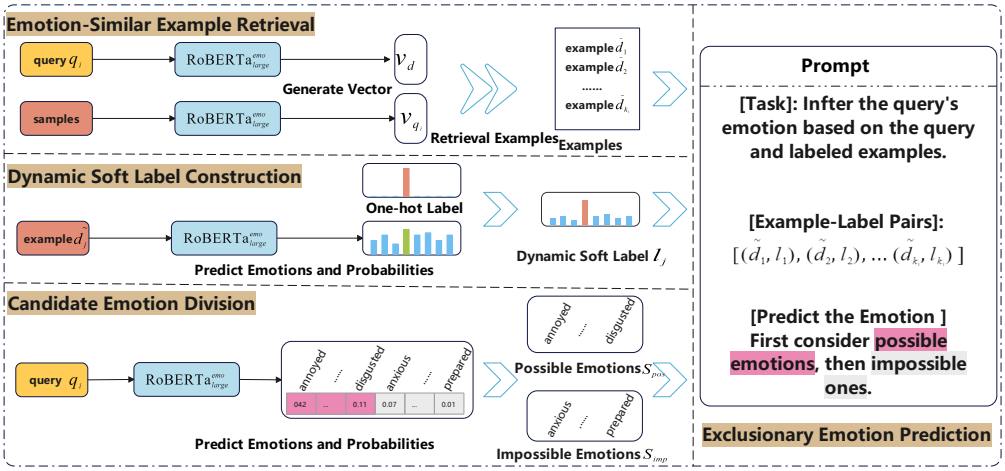

Figure 2: Overview of E-ICL. E-ICL first formulates example-label pairs by retrieving emotion-similar examples and constructing dynamic emotion labels. It then divides emotion categories into possible and impossible emotions, and performs exclusionary emotion prediction based on them. The entire process does not involve additional training, with the assistance of the emotion auxiliary model RoBERTa$^{emo}_{large}$.

methods involve a high cost of manual construction, and for some tasks, the objectives cannot be directly decomposed into sub-process problems. Another part of ICL, i.e., retrieval-based ICL, mitigates this shortcoming by retrieving relevant demonstrations from training datasets. Retrieval-based ICL primarily retrieves demonstrations that are similar to the query in terms of words Rubin et al. (2021); Agrawal et al. (2022); Luo et al. (2023), semantics Li & Qiu (2023); Liu et al. (2021); Yang et al. (2023a); Xiao et al. (2023); Liu et al. (2021), structures Levy et al. (2022), or other relevant aspects Fu et al. (2022); Gonen et al. (2022); Drozdov et al. (2022). Most of these methods rely on the semantics between the query and the demonstrations. Owing to the potential for semantic similarity to result in emotion misunderstanding issues, we propose an emotion-similarity-based retrieval approach and integrate it with an exclusionary emotion prediction mechanism to facilitate more accurate emotion prediction.

3 METHODOLOGY

Task Formulation. We formalize the fine-grained emotion recognition task: given a query Q, the objective is to construct an effective prompt P that guides a large language model to accurately predict the fine-grained emotion category C expressed by Q. Here, Q represents a sentence consisting of M words, and C denotes one of N types of fine-grained emotions.

Overview. The proposed E-ICL is a In-Context Learning-based method, which detects fine-grained emotions through an exclusionary strategy using emotion-similar examples with dynamic soft labels, aided by the emotion-capable auxiliary model RoBERTa$^{emo}_{large}$. As shown inf 2, E-ICL comprises the following four steps: (1) Emotion-Similar Example Retrieval: E-ICL retrieves examples similar in emotion rather than semantics to mitigate emotion misunderstanding of LLMs. (2) Dynamic Soft Label Construction: We dynamically soften the emotion labels to more accurately represent the complex emotions exhibited by the examples. (3) Candidate Emotion Division: We divide the query's emotion into possible and impossible candidates based on the probabilities predicted by the auxiliary model. These candidates are subsequently employed in an exclusionary predictive strategy. (4) Exclusionary Emotion Prediction: Based on the constructed example-label pairs, E-ICL prompts the LLM to predict fine-grained emotion categories through an exclusionary strategy, prioritizing possible emotions over impossible ones. Notably, E-ICL does not involve model training, enabling efficient emotion prediction without consuming substantial computational resources or time.

Emotion Auxiliary Model. We introduce an emotion auxiliary model RoBERTa$^{emo}_{large}$ and leverage its generated emotion vectors and probabilities to enhance large language models (LLMs). Specifically, for an input $Input \in \{D_{test}, D_{train}\}$, we utilize RoBERTa$^{emo}_{large}$ to generate the corresponding vector V and emotion probabilities P.

$$V, P = RoBERTa^{emo}_{large}(Input), \tag{1}$$

where $V \in R^{768}, P \in R^{N_c}$. The generated vector V is used to retrieve emotion-similar examples, while the emotion probabilities P are employed to construct dynamic soft labels and candidate emotions.

Emotion-Similar Example Retrieval. Semantically similar examples can easily lead to emotion misunderstanding issues. To alleviate this problem, we employ an emotion-capable auxiliary model to retrieve emotion-similar examples. Specifically, we map the query $q_i \in D_{test}$ and a sample $s_j \in D_{train}$ into vectors using RoBERTa$^{emo}_{large}$, and calculate their similarity score o_j via the cosine function. We then rank the samples based on the similarity scores, and select the top k_1 highest-scoring samples as the emotion-similar examples d_n.

$$o_j = Cosine(v_{q_i}, v_{d_j}) \tag{2}$$

$$d_n = Top_{k_1}(o_j), n \in [1, k_1] \tag{3}$$

where $v_{q_i}, v_{s_j} \in \mathbb{R}^{768}$ represent the emotion vector representations of the query q_i and the sample s_j, respectively. Top$_{k_1}$ is a ranking function that selects the top k_1 optimal examples. And k_1 is a hyperparameter.

Dynamic Soft Label Construction Emotions are complex, and humans often blend them at varying degrees as needed in daily life Larsen & McGraw (2011); Crivelli & Fridlund (2019); Trampe et al. (2015). Previous In-context learning methods Li & Qiu (2023); Liu et al. (2021) only assign a single deterministic emotion label to the example, which does not align with this human expression habit, leading to an unreasonable label representation. To represent labels more reasonably, we propose a dynamic soft label construction strategy. Specifically, we first use the emotion auxiliary model to predict the emotions e_{d_n} and probabilities p_{d_n} for the example d_n and select the top k_2 emotions with the highest probabilities. For different examples, the predicted emotions and their probabilities vary, endowing the predicted emotions with dynamicity.

$$e_i, p_i = Top_{k_2}(e_{d_n}, p_{d_n}) \tag{4}$$

where $p_i, p_{d_n} \in P, i \in [1, k_2], e_i \in C$. Top$_{k_2}$ is a ranking function that selects the top k_1 optimal emotions. k_2 is a hyperparameter.

Subsequently, we combine the predicted emotions with the ground-truth emotions according to a specific weight α to form dynamic soft labels, where α is a hyperparameter. The specific formula as follows:

$$p_i^{'} = \begin{cases} 1 - \alpha \sum^{k_2} p_j & \text{if } e_i = \text{Ground-Truth Label}, i \in k_2, j \neq i \\ \alpha p_j & \text{Others}, j \in k_2 \end{cases} \tag{5}$$

Candidate Emotion Division. When faced with complex choices, humans tend to first eliminate impossible options and then carefully deliberate among the remaining ones Tversky & Shafir (1992); Dhar (1997); Shafir et al. (1993); Payne et al. (1993). Previous methods Yang et al. (2023a); Xiao et al. (2023) directly predict emotions from multiple emotion categories, such as 41 emotions. This makes the LLMs' choices too abrupt, leading to suboptimal prediction results. In contrast, we mimic the human exclusionary emotion prediction strategy to predict emotions. The first step is to divide the emotion categories into possible and impossible sets. To achieve this goal, we leverage the emotion auxiliary model to predict the query's emotions. We then select the top k_3 emotions with the highest probabilities and consider them as possible emotions, placing them in the set S_{pos}. For the remaining emotions, we consider them as impossible emotions and place them in the set S_{imp}.

$$\widetilde{e}_i = Top_{k_3}(e_{q_i}, p_{q_i}), i \in [1, k_3] \tag{6}$$

$$\widetilde{e}_i \in S_{pos}, S_{imp} \cup S_{pos} = C, S_{imp} \cap S_{pos} = \emptyset \tag{7}$$

where $\widetilde{e}_i, e_i \in C, p_{q_i} \in P$. e_i and p_{q_i} are the emotion categories and probabilities predicted by the auxiliary model, respectively. $\widetilde{e}_i$ represents the possible emotions. Top_{k_3} is a selection function, and k_3 is a hyperparameter.

Exclusionary Emotion Prediction. Based on the above information, we predict fine-grained emotions in an exclusion-based strategy. Specifically, We first construct example-label pairs using emotion-similar examples and dynamic soft labels. Then, we prompt the LLMs to predict emotion types in an exclusion-based manner based on the query and the example-label pairs. In the exclusionary emotion prediction, we prompt the LLMs to first consider the emotion types in the possible emotion set S_{pos} and then consider those in the impossible emotion set S_{imp}. This exclusionary emotion prediction strategy prevents the LLMs from abruptly predicting from all emotions, allowing them to more accurately predict emotions after careful deliberation.

4 EXPERIMENTS

Emotion Auxiliary Model and Datasets. To validate E-ICL, we conduct experiments using different emotion auxiliary models, RoBERTa$^{emo}_{large}$, on various datasets D_{type}, including EDOS Welivita et al. (2021), Empathetic-Dialogues (ED) Rashkin et al. (2019), EmpatheticIntent (EI) Welivita & Pu (2020), and GoEmotions (GE) Demszky et al. (2020). Here, $emo \in \{EI, GE\}$, and $type \in \{EI, GE, ED, EDOS\}$. Note that our goal is to verify the performance of E-ICL without fine-tuning, so the auxiliary model used during inference should not have been fine-tuned on the respective dataset, i.e., $emo \neq type$. Simultaneously, the emotion categories predicted by the auxiliary model do not fully align with those of the datasets, rendering the exclusion strategy inapplicable. To address this issue, we adjust the datasets according to the emotion auxiliary model. For example, for the RoBERTa$^{EI}_{large}$ auxiliary model Welivita & Pu (2020) and the GoEmotions dataset, we first identify the emotion categories they have in common. Then, we select data from GE that falls within these common emotion categories for experimentation. After this adjustment, the available datasets for the RoBERTa$^{EI}_{large}$ auxiliary model are GE, ED, and EDOS, with 19, 32, and 41 emotion categories, respectively. For the RoBERTa$^{GE}_{large}$ [1] auxiliary model, the available datasets are EI, ED, and EDOS, with 19, 17, and 19 emotion categories, respectively.

Evaluation Metrics. We utilize accuracy and macro-F1 for evaluating the methods, following the conventional approach. Accuracy measures the proportion of correctly predicted samples over the total samples. Macro-F1 is the harmonic mean of precision and recall, comprehensively considering both metrics. It accounts for the F1 score of each class and exhibits robustness against class imbalance.

Baselines. E-ICL leverages emotion auxiliary models to enhance the performance of large language models (LLMs) on fine-grained emotion recognition tasks. To validate the proposed method, we first employ the emotion auxiliary models RoBERTa$^{EI}_{large}$ and RoBERTa$^{GE}_{large}$ as baselines. RoBERTa$^{EI}_{large}$ and RoBERTa$^{GE}_{large}$ are RoBERTa models fine-tuned on the EI and GoEmotions emotion datasets, respectively. Secondly, we also select different large language models, namely ChatGPT3-turbo and Claude3-haiku, as baselines. On these LLMs, we construct zero-shot and semantic similarity-based ICL, denoted as Zero-Shot and ICL, respectively.

Implementation Details. Experimental details are provided in Appendix B.

5 RESULTS AND ANALYSIS

5.1 MAIN RESULTS

In Table 1, we list the results when using the emotion auxiliary model RoBERTa$^{EI}_{large}$. Compared to the Zero-Shot and ICL methods, E-ICL exhibits significant advantages, especially on datasets with finer-grained emotions, namely EDOS and ED. This indicates that E-ICL can better perceive and recognize finer-grained emotions in queries. Relative to the emotion auxiliary model RoBERTa$^{EI}_{large}$, E-ICL demonstrates substantial performance improvements across all datasets. Simultaneously, we

[1] https://huggingface.co/mrm8488/roberta-large-bne-finetuned-go_emotions-es

LLM	Models	EDOS		ED		GE	
		Acc	F1	Acc	F1	Acc	F1
-	RoBERTa$^{EI}_{large}$	51.71	52.56	48.96	48.31	24.78	19.64
Claude-haiku	Zero-Shot	25.79	25.10	41.73	36.70	27.65	27.67
	ICL	36.79	38.61	49.47	47.01	36.6	33.04
	E-ICL	**54.23**	**52.78**	**53.98**	**49.2**	**38.05**	**36.80**
ChatGPT-turbo	Zero-Shot	34.6	34.14	36.4	29.82	33.17	29.70
	ICL	39.14	40.04	42.87	41.43	41.37	32.81
	E-ICL	**54.45**	**54.37**	**51.56**	**49.32**	**46.1**	**37.19**

Table 1: Evaluation results on the datasets when using the emotion auxiliary model RoBERTa$^{EI}_{large}$.

LLM	Models	EDOS		ED		EI	
		Acc	F1	Acc	F1	Acc	F1
-	RoBERTa$^{GE}_{large}$	43.25	43.63	40.64	40.07	41.24	41.67
Claude-haiku	Zero-Shot	42.87	37.83	53.22	51.81	53.64	50.57
	ICL	55.73	52.9	61.81	58.88	67.85	64.81
	E-ICL	**62.16**	**57.74**	**62.08**	57.99	66.16	62.04
ChatGPT-turbo	Zero-Shot	54.72	50.66	57.62	55.37	57.81	54.24
	ICL	56.99	54.33	58.18	56.27	61.49	55.28
	E-ICL	**60.4**	**57.00**	**60.85**	**57.65**	**63.05**	**59.91**

Table 2: Evaluation results on the datasets when using the emotion auxiliary model RoBERTa$^{GE}_{large}$.

notice that the emotion auxiliary model exhibits considerable performance discrepancies across different datasets, whereas E-ICL maintains stable performance over multiple datasets. This suggests that E-ICL possesses better robustness compared to the emotion auxiliary.

Table 2 presents the results when using the emotion auxiliary model RoBERTa$^{GE}_{large}$.. On the EDOS dataset, E-ICL significantly outperforms the baselines, demonstrating its advantage in fine-grained emotion recognition. On the ED and EDOS datasets, the ChatGPT-based E-ICL also surpasses the baselines, indicating its effectiveness. However, the Claude-based E-ICL does not exhibit a clear advantage on these two datasets. We find that these two datasets contain some noise, such as the EI dataset having a substantial mixture of English and Chinese characters. When confronted with noisy datasets, the emotion auxiliary model constructs inaccurate dynamic soft labels and eliminated emotion categories. Simultaneously, Claude exhibits lower robustness compared to ChatGPT (as detailed in the Appendix C). Therefore, due to the impact of noise, the Claude-based E-ICL, with lower robustness, exhibits suboptimal performance.

5.2 Analytical Experiments

Ablation Studies. As shown in Figure 3, we conduct ablation studies employing the emotion auxiliary model RoBERTa$^{GE}_{large}$ on the ED, EDOS, and GE datasets. Here, ESE, DSL, and EEP respectively denote emotion-similar examples, dynamic soft labels, and exclusionary emotion prediction, while w/o signifies the absence of the corresponding component. On the EDOS and ED datasets, removing each module results in a certain degree of performance decline for the model. This indicates that emotion-similar examples, dynamic soft labels, and exclusionary emotion prediction all contribute to accurate fine-grained emotion recognition. Simultaneously, on the GE dataset, removing dynamic soft labels (w/o DSL) and the exclusionary strategy (w/o EEP) leads to better performance compared to the complete model. This is primarily due to the presence of substantial noise in this dataset, such as a mixture of Chinese, English, and emoticons, causing the emotion auxiliary model to construct inaccurate dynamic soft labels and eliminated emotions. This largely leads to the suboptimal performance of the model.

Verifying the Contribution of Emotion Auxiliary Model to Dynamic Soft Labels. We verify the impact of the parameter α on the model's performance. The parameter α controls the weight of the emotion probability predicted by the emotion auxiliary model in the dynamic soft labels, where a larger α value indicates a greater contribution from the emotion auxiliary model. To verify its

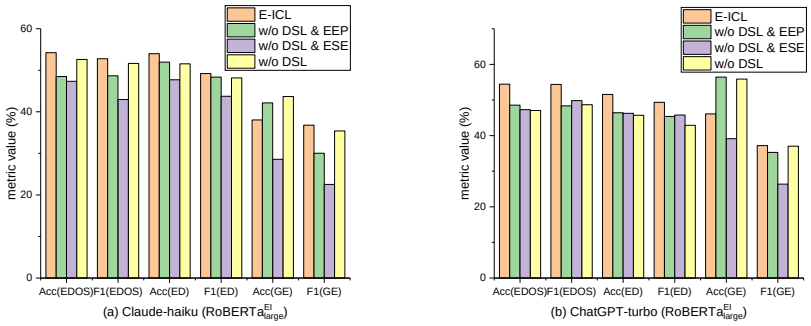

Figure 3: Ablation experiments when using the emotionn auxiliar model RoBERTa$^{EI}_{large}$.

impact, we consider two cases for the emotion auxiliary model: one where its emotion capability is stronger than the employed LLM, and another where it is weaker (Details in Appendix D). We plot the former case in Figure 4(a) and the latter in Figure 4(b). By comparing the two models, we find that when combining with a stronger emotion auxiliary model, the LLM is insensitive to changes in α and consistently exhibits superior performance. Conversely, for a weaker emotion auxiliary model, the LLM should not overly incorporate it, but a moderate combination can bring advantages to the LLM. This is primarily because when the emotion auxiliary model performs poorly, its perception of the query's emotion is inaccurate. Overly considering the emotion auxiliary model introduces more irrelevant noise, leading to a decline in emotion recognition capability.

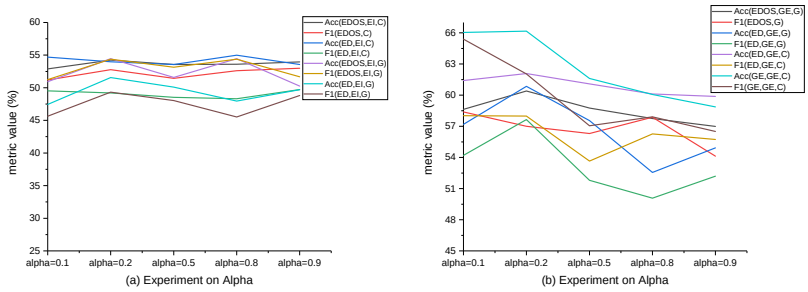

Figure 4: Experimental results of E-ICL based on different α parameters, where elements in the parentheses represent the dataset, the type of emotion auxiliary model, and the type of LLM, respectively. Specifically, C denotes Claude-haiku, while G represents ChatGPT.

Impact of Dynamic soft Label Quantity. We validate the impact of the dynamic soft label number on E-ICL, and the experimental results are shown in Figure 5. Similar to the analysis for the parameter α, we also divide the experiments into two groups based on the emotion capability of the auxiliary models used, from high to low. Figure 5(a) represents the E-ICL based on auxiliary models with stronger emotion capability, while Figure 5(b) represents those with weaker capability. Compared with the two groups of experiments, we find that as the number of soft labels increases, the performance of the former group first decreases and then increases, while the latter group reaches a peak (or starts at a peak) and then decreases. This indicates that a certain number of dynamic soft labels has a promoting effect on emotion prediction for E-ICL. At the same time, excessive labels will introduce noise, which is more obvious for auxiliary models with weaker emotion capability.

Impact of Candidate Emotion Quantity. We conduct experiments with different numbers of candidate emotions $k3$ to validate their impact on the exclusionary emotion prediction strategy. Similarly, we divide the experiments into two groups based on the strength of the emotion auxiliary models. Figure 6(a) is experiments based on stronger emotion auxiliary models, while Figure 6(b) is those based on weaker ones. The experimental results show that considering partial emotion cat-

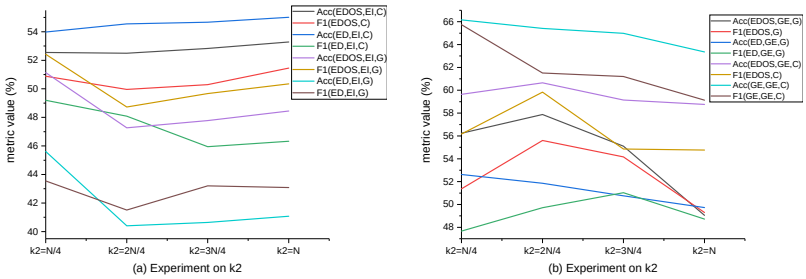

Figure 5: Experimental results of E-ICL based on different k2, where N is the number of emotion categories in the dataset.

egories instead of all categories during the prediction process leads to better performance on most datasets. This demonstrates the effectiveness of the exclusionary emotion prediction strategy. Furthermore, on the EDOS dataset, the E-ICL metrics (Acc: 52.72, F1: 52.14) considering all emotion categories outperform those considering only partial emotions. However, through careful measurement, we find that the Acc and F1 of E-ICL considering only partial emotions are 54.45 and 54.37, respectively, actually outperforming those considering all emotion categories. Therefore, this is due to measurement error. This again demonstrates the effectiveness of the exclusionary emotion prediction strategy.

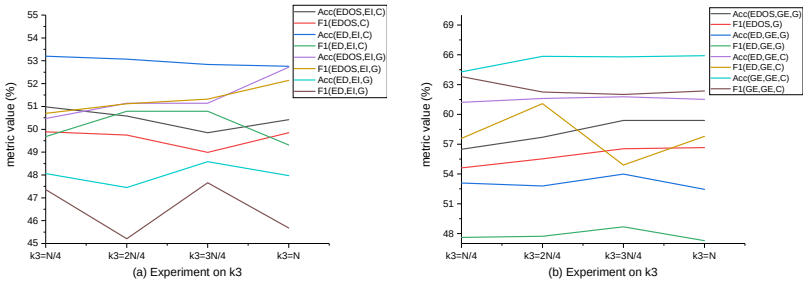

Figure 6: Experimental results of E-ICL based on different k3, where N is the number of emotion categories in the dataset.

More Analysis. To further explore E-ICL, we conduct case studies, as detailed in Appendix E. Concurrently, we analyze the robustness of different LLMs to the introduced noise, as detailed in Appendix C.

6 DISCUSSION

Conclusion. This paper has proposed an Emotional In-Context Learning (E-ICL) approach for fine-grained emotion recognition. From the perspective of prototype theory, E-ICL addresses the shortcomings of ICL in fine-grained emotion recognition: (1) Relying on semantically similar but emotionally inaccurate prototypes for emotion prediction; (2) Being susceptible to interference from irrelevant categories, leading to inaccurate and non-robust judgments during prediction. E-ICL relies on emotionally similar examples with dynamic labels as prototypes for making judgments, avoiding the influence of emotionally inaccurate prototypes. Meanwhile, E-ICL adopts an exclusionary emotion prediction strategy to circumvent interference from irrelevant categories, enabling more accurate and robust emotion prediction. Experimental results and analysis have demonstrated that E-ICL achieves significant advantages in fine-grained recognition tasks without requiring additional computational resources and training time.

Limitations. This paper has the following limitations: (1) Based on prior research Xu et al. (2024); Liu et al. (2021), ICL is likely to conform to prototype theory across a broader range of tasks as well. (2) Semantically similar examples and fixed unique labels may not adequately represent the relationship between examples and labels in certain tasks. (3) The exclusionary prediction strategy benefits ICL's accurate and robust judgments by avoiding interference from irrelevant categories, and it is likely applicable to multi-classification tasks. However, due to resource and time constraints, we do not further explore these limitations in this paper.

Future Work. In the future, we will investigate the following: (1) Explore whether ICL conforms to prototype theory in more tasks; (2) Explore better methods for constructing example-label pairs; (3) Study the applicability of the exclusionary prediction strategy in more tasks.

REFERENCES

Muhammad Abdul-Mageed and Lyle Ungar. EmoNet: Fine-grained emotion detection with gated recurrent neural networks. In Regina Barzilay and Min-Yen Kan (eds.), *ACL*, pp. 718–728, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1067. URL https://aclanthology.org/P17-1067.

Sweta Agrawal, Chunting Zhou, Mike Lewis, Luke Zettlemoyer, and Marjan Ghazvininejad. In-context examples selection for machine translation. *arXiv preprint arXiv:2212.02437*, 2022.

Rafael A Calvo and Sidney D'Mello. Affect detection: An interdisciplinary review of models, methods, and their applications. *IEEE Transactions on affective computing*, 1(1):18–37, 2010.

Carlos Crivelli and Alan J Fridlund. Inside-out: From basic emotions theory to the behavioral ecology view. *Journal of Nonverbal Behavior*, 43(2):161–194, 2019.

Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. Goemotions: A dataset of fine-grained emotions. *arXiv preprint arXiv:2005.00547*, 2020.

Ravi Dhar. Consumer preference for a no-choice option. *Journal of consumer research*, 24(2): 215–231, 1997.

Ning Ding, Xiaobin Wang, Yao Fu, Guangwei Xu, Rui Wang, Pengjun Xie, Ying Shen, Fei Huang, Hai-Tao Zheng, and Rui Zhang. Prototypical representation learning for relation extraction. *arXiv preprint arXiv:2103.11647*, 2021.

Andrew Drozdov, Nathanael Schärli, Ekin Akyürek, Nathan Scales, Xinying Song, Xinyun Chen, Olivier Bousquet, and Denny Zhou. Compositional semantic parsing with large language models. In *The Eleventh International Conference on Learning Representations*, 2022.

Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning. In *The Eleventh International Conference on Learning Representations*, 2022.

Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander Gelbukh. Dialoguegcn: A graph convolutional neural network for emotion recognition in conversation. *arXiv preprint arXiv:1908.11540*, 2019.

Hila Gonen, Srini Iyer, Terra Blevins, Noah A Smith, and Luke Zettlemoyer. Demystifying prompts in language models via perplexity estimation. *arXiv preprint arXiv:2212.04037*, 2022.

James J Gross et al. Emotion regulation: Conceptual and empirical foundations. *Handbook of emotion regulation*, 2:3–20, 2014.

Emitza Guzman and Walid Maalej. How do users like this feature? a fine grained sentiment analysis of app reviews. In *2014 IEEE 22nd international requirements engineering conference (RE)*, pp. 153–162. Ieee, 2014.

James A Hampton. Concepts as prototypes. *Psychology of learning and motivation*, 46:79–113, 2006.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.

Hans Kamp and Barbara Partee. Prototype theory and compositionality. *Cognition*, 57(2):129–191, 1995.

Mehran Kazemi, Najoung Kim, Deepti Bhatia, Xin Xu, and Deepak Ramachandran. Lambada: Backward chaining for automated reasoning in natural language. *arXiv preprint arXiv:2212.13894*, 2022.

Hyunwoo Kim, Byeongchang Kim, and Gunhee Kim. Perspective-taking and pragmatics for generating empathetic responses focused on emotion causes. *arXiv preprint arXiv:2109.08828*, 2021.

Jeff T Larsen and A Peter McGraw. Further evidence for mixed emotions. *Journal of personality and social psychology*, 100(6):1095, 2011.

Itay Levy, Ben Bogin, and Jonathan Berant. Diverse demonstrations improve in-context compositional generalization. *arXiv preprint arXiv:2212.06800*, 2022.

Qintong Li, Hongshen Chen, Zhaochun Ren, Pengjie Ren, Zhaopeng Tu, and Zhumin Chen. Empdg: Multiresolution interactive empathetic dialogue generation. *arXiv:abs/1911.08698*, 2020.

Qintong Li, Piji Li, Zhaochun Ren, Pengjie Ren, and Zhumin Chen. Knowledge bridging for empathetic dialogue generation. In *AAAI*, 2022.

Xiaonan Li and Xipeng Qiu. Mot: Pre-thinking and recalling enable chatgpt to self-improve with memory-of-thoughts. *arXiv preprint arXiv:2305.05181*, 2023.

Jasy Suet Yan Liew and Howard R Turtle. Exploring fine-grained emotion detection in tweets. In *Proceedings of the NAACL student research workshop*, pp. 73–80, 2016.

Zhaojiang Lin, Andrea Madotto, Jamin Shin, Peng Xu, and Pascale Fung. Moel: Mixture of empathetic listeners. In *EMNLP-IJCNLP*, pp. 121–132, 2019.

Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*, 2021.

Man Luo, Xin Xu, Zhuyun Dai, Panupong Pasupat, Mehran Kazemi, Chitta Baral, Vaiva Imbrasaite, and Vincent Y Zhao. Dr. icl: Demonstration-retrieved in-context learning. *arXiv preprint arXiv:2305.14128*, 2023.

Navonil Majumder, Soujanya Poria, Devamanyu Hazarika, Rada Mihalcea, Alexander Gelbukh, and Erik Cambria. Dialoguernn: An attentive rnn for emotion detection in conversations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 6818–6825, 2019.

Navonil Majumder, Pengfei Hong, Shanshan Peng, Jiankun Lu, Deepanway Ghosal, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. Mime: Mimicking emotions for empathetic response generation. In *EMNLP*, pp. 8968–8979, 2020.

Agnes Moors, Phoebe C Ellsworth, Klaus R Scherer, and Nico H Frijda. Appraisal theories of emotion: State of the art and future development. *Emotion review*, 5(2):119–124, 2013.

Meike Nauta, Ron Van Bree, and Christin Seifert. Neural prototype trees for interpretable fine-grained image recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14933–14943, 2021.

John W Payne, James R Bettman, and Eric J Johnson. *The adaptive decision maker*. Cambridge university press, 1993.

Wei Peng, Yue Hu, Luxi Xing, Yuqiang Xie, Yajing Sun, and Yunpeng Li. Control globally, understand locally: A global-to-local hierarchical graph network for emotional support conversation. *arXiv preprint arXiv:2204.12749*, 2022.

Rosalind W Picard. *Affective computing*. MIT press, 2000.

Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. A review of affective computing: From unimodal analysis to multimodal fusion. *Information fusion*, 37:98–125, 2017.

Yushan Qian, Wei-Nan Zhang, and Ting Liu. Harnessing the power of large language models for empathetic response generation: Empirical investigations and improvements. *arXiv preprint arXiv:2310.05140*, 2023.

Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.

Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *ACL*, pp. 5370–5381, 2019.

Eleanor Rosch. Principles of categorization. In *Cognition and categorization*, pp. 27–48. Routledge, 1978.

Ohad Rubin, Jonathan Herzig, and Jonathan Berant. Learning to retrieve prompts for in-context learning. *arXiv preprint arXiv:2112.08633*, 2021.

Sahand Sabour, Chujie Zheng, and Minlie Huang. Cem: Commonsense-aware empathetic response generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Virginia, USA, 2022. AAAI Press.

Tulika Saha, Saraansh Chopra, Sriparna Saha, Pushpak Bhattacharyya, and Pankaj Kumar. A large-scale dataset for motivational dialogue system: An application of natural language generation to mental health. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2021.

Tulika Saha, Vaibhav Gakhreja, Anindya Sundar Das, Souhitya Chakraborty, and Sriparna Saha. Towards motivational and empathetic response generation in online mental health support. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pp. 2650–2656, 2022.

Kristina Schaaff, Caroline Reinig, and Tim Schlippe. Exploring chatgpt's empathic abilities. In *2023 11th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pp. 1–8. IEEE, 2023.

Hendrik Schuff, Jeremy Barnes, Julian Mohme, Sebastian Padó, and Roman Klinger. Annotation, modelling and analysis of fine-grained emotions on a stance and sentiment detection corpus. In *Proceedings of the 8th workshop on computational approaches to subjectivity, sentiment and social media analysis*, pp. 13–23, 2017.

Eldar Shafir, Itamar Simonson, and Amos Tversky. Reason-based choice. *Cognition*, 49(1-2):11–36, 1993.

Matthias Siemer, Iris Mauss, and James J Gross. Same situation–different emotions: how appraisals shape our emotions. *Emotion*, 7(3):592, 2007.

Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.

Debra Trampe, Jordi Quoidbach, and Maxime Taquet. Emotions in everyday life. *PloS one*, 10(12): e0145450, 2015.

Quan Tu, Yanran Li, Jianwei Cui, Bin Wang, Ji-Rong Wen, and Rui Yan. Misc: a mixed strategy-aware model integrating comet for emotional support conversation. *arXiv preprint arXiv:2203.13560*, 2022.

Amos Tversky and Eldar Shafir. The disjunction effect in choice under uncertainty. *Psychological science*, 3(5):305–310, 1992.

Yufeng Wang, Chen Chao, Yang Zhou, Wang Shuhui, and Liao Xiangwen. Ctsm: Combining trait and state emotions for empathetic response model. *arXiv preprint arXiv:2403.15516*, 2024.

Zhaoxia Wang, Chee Seng Chong, Landy Lan, Yinping Yang, Seng Beng Ho, and Joo Chuan Tong. Fine-grained sentiment analysis of social media with emotion sensing. In *2016 Future Technologies Conference (FTC)*, pp. 1361–1364, 2016. doi: 10.1109/FTC.2016.7821783.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

Anuradha Welivita and Pearl Pu. A taxonomy of empathetic response intents in human social conversations. *arXiv preprint arXiv:2012.04080*, 2020.

Anuradha Welivita, Yubo Xie, and Pearl Pu. A large-scale dataset for empathetic response generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 1251–1264, 2021.

Chaojun Xiao, Zhengyan Zhang, Xu Han, Chi-Min Chan, Yankai Lin, Zhiyuan Liu, Xiangyang Li, Zhonghua Li, Zhao Cao, and Maosong Sun. Plug-and-play document modules for pre-trained models. *arXiv preprint arXiv:2305.17660*, 2023.

Yubo Xie, Ekaterina Svikhnushina, and Pearl Pu. A multi-turn emotionally engaging dialog model. *arXiv preprint arXiv:1908.07816*, 2019.

Xin Xu, Yue Liu, Panupong Pasupat, Mehran Kazemi, et al. In-context learning with retrieved demonstrations for language models: A survey. *arXiv preprint arXiv:2401.11624*, 2024.

Linyi Yang, Shuibai Zhang, Zhuohao Yu, Guangsheng Bao, Yidong Wang, Jindong Wang, Ruochen Xu, Wei Ye, Xing Xie, Weizhu Chen, et al. Supervised knowledge makes large language models better in-context learners. *arXiv preprint arXiv:2312.15918*, 2023a.

Zhou Yang, Zhaochun Ren, Wang Yufeng, Xiaofei Zhu, Zhihao Chen, Tiecheng Cai, Wu Yunbing, Yisong Su, Sibo Ju, and Xiangwen Liao. Exploiting emotion-semantic correlations for empathetic response generation. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023b. URL https://openreview.net/forum?id=ilCMZV0Qdl.

Zhou Yang, Zhaochun Ren, Wang Yufeng, Shizhong Peng, Haizhou Sun, Xiaofei Zhu, and Xiangwen Liao. Enhancing empathetic response generation by augmenting llms with small-scale empathetic models. *arXiv preprint arXiv:2402.11801*, 2024.

Weixiang Zhao, Yanyan Zhao, Xin Lu, and Bing Qin. Don't lose yourself! empathetic response generation via explicit self-other awareness. *arXiv preprint arXiv:2210.03884*, 2022.

Weixiang Zhao, Yanyan Zhao, Xin Lu, Shilong Wang, Yanpeng Tong, and Bing Qin. Is chatgpt equipped with emotional dialogue capabilities? *arXiv preprint arXiv:2304.09582*, 2023.

A Pilot Experiments

We conduct experiments on four fine-grained emotion recognition datasets: EDOS, Empathetic-Dialogues, EmpatheticIntent, and GoEmotions. Taking Empathetic-Dialogues as an example, we first construct eight sets of examples, with each set containing five example-label pairs. Second, we treat the constructed example-label pairs as demonstrations and map them into vectors using RoBERTa$_{large}$. Subsequently, for each demonstration, we select 1216 queries based on similarity scores. We then assemble the demonstrations and queries as inputs to prompt the LLMs for emotion prediction. To eliminate interference from different LLMs, we perform experiments on both ChatGPT-turbo and Claude-haiku.

B Implementation Details

In our experiments, we employ two emotion auxiliary models, RoBERTa$^{EI}_{large}$ and RoBERTa$^{GE}_{large}$, with the former being used for validation on the GE, ED, and EDOS datasets, while the latter is used for EI, ED, and EDOS. During the construction of the instance-label pairs, the example number is

set to $k_1 = 5$. Meanwhile, the fusion weight for the soft labels is $\alpha = 0.2$. Additionally, the number of soft labels k_2 and the number of impossible emotions k_3 are influenced by variations in data, emotion auxiliary models, and LLMs. Consequently, we provide a detailed analysis and discussion of these factors in section 5.2.

C Appendix: Robustness Analysis of LLMs

To validate the robustness of LLMs to E-ICL, we conduct the following experiments. First, we select Model RoBERTa$^{GE}_{large}$ as the emotion auxiliary model. Since this model performs relatively poorly on the respective EDOS, ED, and EI datasets, it will introduce more noise, which is beneficial for robustness experiments. Next, we select different numbers of candidate emotions to validate the robustness of Claude-haiku and ChatGPT-turbo. The experimental results are shown in Figure 3. The x-axis represents the number of candidate emotions, and the y-axis represents the metric values. As shown in Figure 7, as the number of candidate emotions (and noise) increases, the metric values of Claude-haiku fluctuate significantly, while those of ChatGPT-turbo remain stable within a certain range. This indicates that ChatGPT-turbo is more robust to the noise introduced by E-ICL.

D Appendix: Grouped Analysis of Emotion Auxiliary Models

As shown in Tables 3 and 4, the emotion auxiliary models exhibit different performance across different datasets. Ignoring these differences and directly analyzing the experiments would lead to unreliable results. To investigate their impact, we divide the emotion auxiliary models into two groups: (a) those that significantly outperform LLMs, and (b) those that underperform LLMs. Specifically, we find that when using RoBERTa$^{EI}_{large}$ on the EDOS and ED datasets, its performance is significantly better than Claude-haiku and ChatGPT-turbo, while on the GE dataset, it underperforms compared to them. Therefore, we categorize the experiments based on RoBERTa$^{EI}_{large}$ and conducted on the EDOS and ED datasets as the (a) group experiments, while the experiments on the ED dataset are categorized as the (b) group experiments. Simultaneously, we adopt the same approach to divide the experiments based on RoBERTa$^{GE}_{large}$. Since RoBERTa$^{GE}_{large}$ does not significantly outperform LLMs on the EDOS, ED, and GE datasets, we categorize its experiments as the (b) group. When conducting parameter analysis experiments, due to the performance differences between the emotion auxiliary models and LLMs, the (a) group experiments and (b) group experiments exhibit different characteristics. This division of experiments better shows the impact of the emotion auxiliary models.

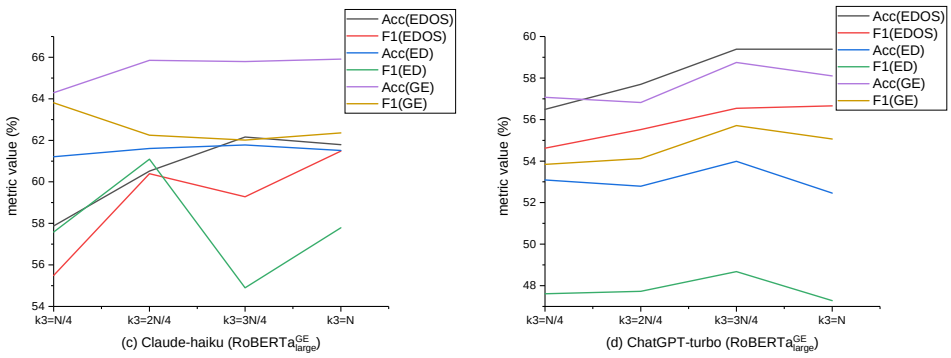

Figure 7: Experimental results of RoBERTa$^{GE}_{large}$ based E-ICL on different k3, where N is the number of emotion categories in the dataset.

Comparison	EDOS		ED		GE	
	Acc	Macro-F1	Acc	Macro-F1	Acc	Macro-F1
RoBERTa$^{EI}_{large}$ vs. Claude	25.92	27.46	7.23	11.61	-2.87	-8.03
RoBERTa$^{EI}_{large}$ vs. ChatGPT	17.11	18.42	12.56	18.49	-8.39	-10.06

Table 3: Comparisons between RoBERTa$^{EI}_{large}$ and LLMs on different datasets, where positive values indicate RoBERTa$^{EI}_{large}$ outperforming LLMs, and negative values indicate the opposite.

Comparison	EDOS		ED		EI	
	Acc	Macro-F1	Acc	Macro-F1	Acc	Macro-F1
RoBERTa$^{GE}_{large}$ vs. Claude	0.38	5.8	-12.58	-11.74	-12.4	-8.9
RoBERTa$^{GE}_{large}$ vs. ChatGPT	-11.47	-7.02	-16.98	-15.3	-16.57	-12.57

Table 4: Comparison between RoBERTa$^{GE}_{large}$ and LLMs on different datasets, where positive values indicate RoBERTa$^{GE}_{large}$ outperforming LLMs, and negative values indicate the opposite.

E Appendix: Case Study

To demonstrate the advantages of E-ICL, we conduct a case study. The analysis results are shown in Table 5. The query of the 1st case expresses the emotion of "caring." The Zero-Shot method cannot accurately perceive this fine-grained emotion. In-context learning (ICL) predicts the query's emotion by retrieving and understanding semantically similar examples. However, the emotions of the semantically similar examples are diverse, such as "agreeing," "caring," and "grateful." Due to the difficulty in distinguishing among various emotions, ICL fails to accurately judge the emotion of the query, leading to an incorrect prediction of "encouraging." In contrast, E-ICL predicts the emotion of the query by retrieving and understanding examples with similar emotions, accurately predicting the query's emotion as "caring."

In the second case, the query expresses the emotion of "jealous." Similarly, the Zero-Shot method cannot accurately perceive this subtle emotion type. In the ICL method, the retrieved examples semantically similar to the query have diverse emotions, making it difficult for the LLM to accurately determine the emotion type based on these examples. In contrast, E-ICL retrieves three examples with similar emotions, enabling the LLM to make a more accurate prediction combined with these examples.

Query	Hi , Tommy . I'm delivering a gift to Addie that's going to help get her back on track. **Emotion**: Caring		
Methods	E-ICL	ICL	Zero-Shot
Example 1	So what brought you back here ?¡/s¿When Mr Winters died they didn't have a replacement. I decided I'm going to rescue these poor kids. **Emotion**: Caring	It's your turn, mum .¡/s¿I know! **Emotion**: Agreeing	–
Example 2	My grandmother's not doing so wel, so I took a year off from school to help her out. **Emotion**: Caring	Hey, there, Josey !¡/s¿We're going to Santo Rio! **Emotion**: Excited	–
Example 3	Pancho has a good heart. He feeds his little pet. **Emotion**: Caring	Good morning, Mr. Stark .¡/s¿I brought you some homemade cookies. **Emotion**: Caring	–
Example 4	We came to the prisoners. **Emotion**: Caring	I have something for you. You'll be heard at Salamanca University . ¡/s¿ In a week! **Emotion**: Anticipating	–
Example 5	We built Graciela's casitas for abandoned women and children who needed a place to stay. **Emotion**: Sentimental	There you go, free sweets up for grabs. All you've got to do is get them out of the tube .¡/s¿We placed everything needed within reach. **Emotion**: Grateful	–
Prediction	Caring	Encouraging	Hopeful
Query	I think she wants all the women around her to look fat. **Emotion**: Jealous		
Methods	E-ICL	ICL	Zero-Shot
Example 1	Can you imagine that ? All the women are around me in my office all day long, she's jealous over some foreign country I've never been to before. **Emotion**: Jealous	Everybody , listen ! Mr Anderson wants his team to play us .¡/s¿A Japan-America All-Star game! **Emotion**: Anticipating	–
Example 2	Kimura . Jealousy makes me feel much younger. **Emotion**: Jealous	Natasha tries to get me out here once a week. **Emotion**: Annoyed	–
Example 3	Do you think my body is beautiful ?¡/s¿I hate beautiful bodies . **Emotion**: Disgusted	Sanga ... Come and see her son do this . **Emotion**: Proud	–
Example 4	See those beauties ?¡/s¿Know them ? Nope .¡/s¿I wish! The one in red is hot. **Emotion**: Hopeful	According to my source, he's your gigolo .¡/s¿I thought that men were the only ones who wanted to own women. But even young women have their personal toys. **Emotion**: Jealous	–
Example 5	They look like they're daring each other to move in. I hate that when a guy comes up to hit on you while his friends watch. **Emotion**: Jealous	He says he's waitin' for the presents. **Emotion**: Anticipating	–
Prediction	Jealous	Annoyed	Disgusted

Table 5: Case study of E-ICL and Benchmarks.