

Investigating the Learning Behaviour of In-Context Learning: A Comparison with Supervised Learning

Xindi Wang^{a, b;*}, Yufei Wang^{c;**}, Can Xu^d, Xiubo Geng^d, Bowen Zhang^d, Chongyang Tao^d,
Frank Rudzicz^{b, e, f;***}, Robert E. Mercer^{a, b;****} and Daxin Jiang^{d;*****}

^aUniversity of Western Ontario, Canada

^bVector Institute for Artificial Intelligence, Canada

^cMacquarie University, Sydney, Australia

^dMicrosoft Corporation, Beijing, China

^eUniversity of Toronto, Canada

^fDalhousie University, Canada

ORCID ID: Xindi Wang <https://orcid.org/0000-0002-6548-3547>,

Frank Rudzicz <https://orcid.org/0000-0002-1139-3423>, Robert E. Mercer <https://orcid.org/0000-0002-0080-715X>

Abstract. Large language models (LLMs) have shown remarkable capacity for in-context learning (ICL), where learning a new task from just a few training examples is done without being explicitly pre-trained. However, despite the success of LLMs, there has been little understanding of how ICL learns the knowledge from the given prompts. In this paper, to make progress toward understanding the learning behaviour of ICL, we train the same LLMs with the same demonstration examples via ICL and supervised learning (SL), respectively, and investigate their performance under label perturbations (i.e., noisy labels and label imbalance) on a range of classification tasks. First, via extensive experiments, we find that gold labels have significant impacts on the downstream in-context performance, especially for large language models; however, imbalanced labels matter little to ICL across all model sizes. Second, when comparing with SL, we show empirically that ICL is less sensitive to label perturbations than SL, and ICL gradually attains comparable performance to SL as the model size increases.

1 Introduction

Recent advances in large-scale pre-trained language models (LLMs), such as GPT-3 [3], have led to an interesting emergent learning paradigm called *in-context learning* (ICL). In the ICL paradigm, given a prompt that includes a list of few-shot training input-output data and a test input at the end, LLMs directly make a prediction conditioning on the prompt without any updates to their model parameters. This is in contrast with current standards in fine-tuning [9, 32], where model parameters are updated according to the gradients of training losses. Compared to supervised fine-tuning (supervised learning), ICL allows users to directly manipulate LLMs with only language-based prompts and more modest computational resource requirements.

However, despite the advantages of ICL, it is still unclear how ICL learns knowledge from the given prompts without updating its model parameters. Preliminary research [1, 11] compared ICL with simple machine learning models, such as logistic regression and shallow neural networks. In this paper, we take a further step and investigate learning behaviour differences between ICL and supervised learning (SL). Specifically, we train three LLMs with the same training data via *in-context learning* and *supervised learning* separately and analyze their generated outputs. While SL is a well-established approach that uses labelled data to train models to make accurate predictions, ICL takes a different approach by leveraging the context of the text to learn from unlabeled data in order to improve the accuracy of the predictions. By comparing the performance of ICL and SL, we gain insights into the effectiveness and weaknesses of each approach. In addition, previous work on ICL has hinged upon clean and balanced data; however, in practice, these conditions are incredibly difficult and expensive to meet. In situations with unclean or imbalanced data, ICL may provide a more cost-effective approach to learning from limited available data, and comparing it with SL can help understand its potential advantages and limitations. Inspired by previous work which uses perturbed data to investigate properties of deep neural networks [14], we apply *label perturbations* (i.e., incorrectly annotated labels and imbalanced distributed labels) to the above training data and observe the corresponding performance changes in both types of learning paradigms. This differs from previous ICL research [45] which only used balanced training data with high-quality annotations.

In this paper, we investigate experimentally the learning behaviour of in-context learning and compare it with supervised learning under label perturbations, as shown in Figure 1. We compare model performance over six text classification datasets with different model sizes, and we evaluate the sensitivity of supervised learning and in-context learning on perturbed and clean data. We also provide a complementary study as to why the sensitivity of ICL differs with two label perturbation settings by calculating the attention scores on the labels. We provide the following major empirical findings on ICL un-

* Work done during an internship at Microsoft STCA.

** Work done during an internship at Microsoft STCA.

*** Senior author.

**** Senior author.

***** Corresponding Author. Email: djjiang@microsoft.com.

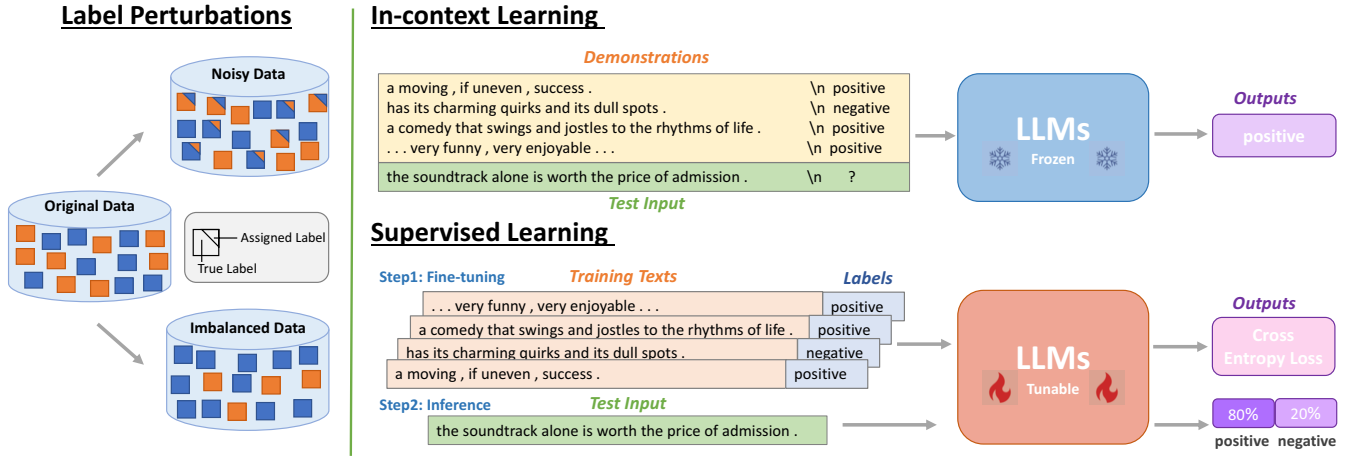


Figure 1. A overview of our empirical studies. There are three main steps in our experiments. First, construct perturbed data from the original data. Second, we perform in-context learning: the demonstration examples consist of k input-label pairs from the training data ($k = 4$ in this figure). Lastly, supervised learning: first fine-tune the LM using the same data as demonstrations in ICL and then use the fine-tuned LM for inference.

der label perturbations: (1) gold labels are important to ICL and large language models are more sensitive to corrupted labels than smaller models; (2) label imbalance matters little to ICL across all model sizes. In comparing the two learning paradigms (i.e., ICL and SL), we observe the following: (1) in-context learning is less sensitive to label perturbations than supervised learning; (2) in-context learning gradually attains comparable performance to supervised learning as the model size increases. The contributions of this paper are:

1. We provide a new way of understanding the learning behaviour of ICL by using perturbed labels.
2. We propose an explanation of the learning behaviour of ICL by using attention scores.
3. We are the first to perform a systematical comparison of the learning behaviours of ICL and SL under label perturbations.

2 Terminology

2.1 Problem Formulation

In this paper, we explore in-context learning in downstream classification task \mathcal{T} which includes a discrete output label set $\mathcal{Y} = \{y_1, y_2, \dots, y_L\}$ and corresponding training data $\mathcal{D}_{\mathcal{T}} = \{(x_i, y_i)\}, i = 1, 2, \dots, N$, where x_i is the input text, $y_i \in \mathcal{Y}$, and N is the number of input-output pairs in $\mathcal{D}_{\mathcal{T}}$. Our goal is to learn a model \mathcal{M} that predicts the correct y_j using x_i as input.

Supervised Learning In the traditional supervised learning setting, given a training instance (x_i, y_j) , we perform the parameter update for \mathcal{M} based on the training loss, as follows:

$$\mathcal{L}(\mathcal{M}(x_i), \mathbf{y}_j, \boldsymbol{\theta}), \quad (1)$$

where \mathbf{y}_j is the one-hot representation for y_j and $\boldsymbol{\theta}$ are the parameters of \mathcal{M} .

In-context Learning Instead of training with gradient descent, *In-Context Learning* directly estimates y_i based on the input text x_i and k demonstration examples $\mathcal{C} = (x_1^c, y_1^c), \dots, (x_k^c, y_k^c)$. The final prediction y_i is chosen from \mathcal{Y} , as follows:

$$y_i = \underset{j}{\operatorname{argmax}} \mathcal{M}(y_j | \mathcal{C}, x_i), \quad (2)$$

where \mathcal{C} is represented by a concatenation of input-label pairs using pre-defined templates. We refer the reader to Table 1 for the detailed templates that are used for each dataset.

2.2 Learning with Label Perturbations

Previous studies on *in-context learning* mostly focus on the clean and high-quality data settings where the training data has high-quality annotations and roughly equal distributions over different labels in \mathcal{Y} . However, in many real-world scenarios, meeting such requirements could be challenging. Furthermore, as shown in [14], perturbed data could also be helpful in understanding the behavior of many neural networks. Thus, in this paper, we investigate how *in-context learning* reacts to perturbed training data and compare such behavior with *supervised learning*. Specifically, we consider two types of perturbations that simulate different levels of noise and degrees of class imbalance in real-world situations.

Noisy Labels In learning with noisy labels, we inject noise to form a label-corrupted training set $\tilde{\mathcal{D}}_{\mathcal{T}} = \{(x_1, \tilde{y}_1), (x_2, \tilde{y}_2), \dots, (x_N, \tilde{y}_N)\}$, that is obtained from a noisy joint distribution over $\mathcal{X} \times \mathcal{Y}$, where $\tilde{\mathcal{D}}$ contains a combination of clean samples (whose original label is clean $y_i = \tilde{y}_i$) and noisy samples (whose original label is corrupted $y_i \neq \tilde{y}_i$). Our goal is to learn a competitive classifier by training on corrupted data $\tilde{\mathcal{D}}_{\mathcal{T}}$. We consider uniform (or symmetric) noise with noise levels ϵ .

The uniform noise flips a label from its true class to any other class with equal probability, p_{uni} , that can be defined as:

$$p_{uni}(\tilde{y} = j | y = i) = \begin{cases} 1 - \epsilon, & \text{for } i = j \\ \frac{\epsilon}{L-1}, & \text{for any } i \neq j \end{cases}, \quad (3)$$

where $p_{uni}(\tilde{y} = j | y = i)$ represents the probability of changing the clean label i to the noisy label j , and ϵ is the noise level.

Label Imbalance In practice, skewed data distributions arise in many datasets. In learning with label imbalance, the degree of class imbalance can be represented by the imbalance ratio. Given a balanced data distribution \mathcal{D} , let r denote the ratio of class imbalance,

Dataset	Input	Label	Format	
			Model	Template
MR	a moving , if uneven , success .	positive, negative	GPT2-Large	Input_Label_
			GPT2-XL	Input_Label_
			GPT-J	Input \n Label \n\n\n
SST2	sentence: well worth revisiting as many times	positive, negative	GPT2-Large	Input_Label_
			GPT2-XL	Input_Label_
			GPT-J	Input \n Label \n\n\n
RTE	Alleged terrorists today, killed Dolores Hinostrroza, the mayor of Mulqui district, shooting her five times. implies: The mayor of Mulqui district was murdered with a firearm. True or False?	True, False	GPT2-Large	Input_Label_
			GPT2-XL	Input_Label_
			GPT-J	Input \n Label \n\n\n
CB	B: Right, you know, like In packaging A: Yeah. B: and, uh, you know, just goodness. A: Yeah, I don't think they do the packaging at this plant, question: they do the packaging at this plant. Entailment, contradiction, or neutral?	entailment, contradiction, neutral	GPT2-Large	Input_Label_
			GPT2-XL	Input_Label_
			GPT-J	Input \n Label \n\n\n
AG-NEWS	Toshiba is taking back bad memory ZDNet's survey of IT professionals in October kept upgrading hardware at number two on the businesses radar throughout the year.	Business, Science, Sports, World	GPT2-Large	Input_Label_
			GPT2-XL	Input_Label_
			GPT-J	Input \n Label \n\n\n
TREC	What actor came to dinner in Guess Who 's Coming to Dinner ?	Abbreviation, Description, Entity, Person, Location, Number	GPT2-Large	Input_Label_
			GPT2-XL	Input_Label_
			GPT-J	Input \n Label \n\n\n

Table 1. A list of prompt templates we used for ICL. We show one example per task in the demonstration for illustration purposes. (Note: _ represents space, and \n represents new line.)

\mathcal{D}^r is the imbalanced data distribution for ratio r . That is, r is the probability of the rarest class over that of the most frequent class:

$$r = \frac{\min_{l \in \mathcal{Y}} \mathcal{D}(y=l)}{\max_{l \in \mathcal{Y}} \mathcal{D}(y=l)}, \quad (4)$$

where $0 \leq r \leq 1$. The dataset is balanced if $r = 1$ and the dataset is heavily long-tailed if r is small.

2.3 Sensitivity to Label Perturbation

Here, sensitivity is the degree to which the downstream classification performance changes when the model is subject to a fixed amount of label perturbation. Following Yoo *et al.* [41], we measure the sensitivity using a linear regression analysis on the model performance against the perturbed rates:

$$f_{\mathcal{M}} = \beta_0 + \beta_1 \mathcal{P}, \quad (5)$$

where $f_{\mathcal{M}}$ is the performance of language model \mathcal{M} , and \mathcal{P} is the perturbation rate. The scalar value β_1 is interpreted as the sensitivity measure.

In general, the performance of a model drops with perturbation. The coefficient of the linear regression is usually negative, and a lower absolute coefficient value indicates weak sensitivity. The intercept β_0 represents the performance of gold labels.

3 Experiments

3.1 Datasets

We conduct experiments on six datasets: binary sentiment classification with Movie Review (MR) [29] and SST-2 [34]; textual entailment using RTE [7] and CommitmentBank (CB) [8]; topic classification using AGNews [43]; and question classification using TREC

[19, 13]. Table 1 shows the default prompt format used for all datasets. For GPT2-Large and GPT2-XL, we separate the input and the label and each demonstration example with a space. For GPT-J, we separate the input and the label with a new line and each demonstration example with three new lines. The statistics of the datasets as well as the corruption rate and imbalance ratio for each dataset are summarized in Table 2.

We perturb the datasets in two ways:

- Noisy labels: we corrupt labels by various amounts from 0% to 100%, with steps of 25%, with uniform noise.
- Label imbalance: we construct label distributions with three levels of imbalance ratios: low (66% to 100%), medium (33% to 66%), and high (0% to 33%).

Dataset	Classes	Number of Testing	Noise Level (%)	Imbalance Ratio (%)
MR	2	1066	100/ 75/ 50/ 25/ 0	100/ 45/ 7
SST-2	2	872	100/ 75/ 50/ 25/ 0	100/ 45/ 7
RTE	2	277	100/ 75/ 50/ 25/ 0	100/ 45/ 7
CB	3	56	100/ 75/ 50/ 25/ 0	83/ 43/ 10
AGNews	4	7600	100/ 75/ 50/ 25/ 0	100/ 50/ 10
TREC	6	1091	100/ 75/ 50/ 25/ 0	67/ 33/ 9

Table 2. Details of the text classification datasets.

3.2 Implementation Details

We use 16 examples as demonstrations by default for both SL and ICL. We use 5 different random seeds and run all experiments 5 times. For SL, we first do grid search for hyper-parameter optimization on each dataset using 5 different seeds and then fine-tune the model with the optimum values of hyper-parameters per seed. We

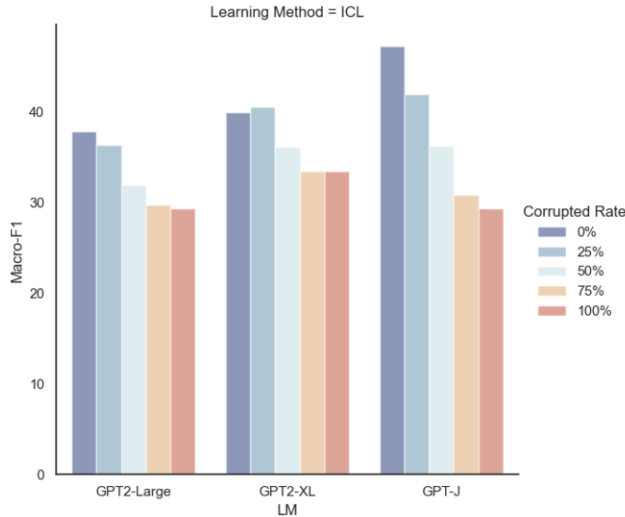


Figure 2. Results with varying ratio of correct labels for in-context learning over three different models.

experiment with three language models, GPT2-Large [32], GPT2-XL [32], and GPT-J [35], ranging from 774 million to 6.7 billion parameters, all being decoder-only models. Evaluating on a comparable set of models with different sizes allows us to further investigate whether the observations depend on the model size. We report Macro-F1 scores for all tasks and compute the average score over seeds for each dataset. The code for reproducing our experiments is available at https://github.com/xdwang0726/ICL_LL.

4 Empirical Findings

We intensively experiment to answer the following:

1. Does the ratio of correct labels matter to the performance of in-context learning?
2. Is in-context learning less sensitive to noisy labels compared to supervised learning?
3. Does the imbalance ratio affect in-context learning performance?
4. Is in-context learning more sensitive to label imbalance compared to supervised learning?

Tables 3 and 4 report the main results for in-context learning and supervised learning performance across three different models under noisy label and label imbalance settings, respectively¹.

4.1 How do ICL and SL Learn with Noisy Labels?

Ground-truth Labels matter a lot to ICL Min et al. [27] claimed that replacing ground-truth labels with incorrect labels marginally affects the overall performance on specific datasets. In contrast, our work conducts extensive experiments on the impact of gold labels, and it is noticeable that the ability of ICL emerges differently depending on corruption rates, as shown in Figure 2. Model performance is sensitive to the ratio of correct labels in the demonstrations. In fact, using all corrupted labels always significantly under-performs demonstration with all gold labels across all three models. Meanwhile, as the model size increases, the performance gap between various corruption rates also increases, which indicates that the large

language models are more sensitive to the corrupted labels. For instance, the performance gap between all gold and all corrupted labels for GPT2-Large is 8.5%, and for GPT-J is 17.9%.

Correct labels matter a lot to the model performance but, in rare cases, model performance is not guaranteed with lower corruption rates. An inflection point might exist where a certain number of corrupted labels helps *improve* model performance. For instance, 75% correct labels performs slightly better than using all gold labels for GPT2-XL.

Compared with SL Overall, as the corrupted rate increases, the performance of supervised learning declines dramatically, and the performance of ICL drops slowly for GPT2-Large and GPT2-XL but decreases as rapidly as SL for GPT-J, as shown in Figure 3. Looking at the sensitivity measures in Table 5, we see that ICL is less sensitive than SL across all three models. In addition, ICL gradually outperforms SL as the model size increases. SL performance drops up to 19% after injecting incorrect labels, while ICL performance drops 11%. It is also noticeable that, as more incorrect labels are introduced, the performance of ICL declines gradually at first and then becomes constant, while SL steadily declines. ICL gradually attains performance comparable to that of SL as the model size increases.

We are also interested in how various learning paradigms perform with different classification types. We divide the six classification tasks into two types, binary classification and multi-class classification. As shown in Table 3², there is no large gap between the performance of SL and ICL in binary classification. Specifically, in binary classification with GPT2-large and GPT-J, ICL is slightly better with lower corruption rates and SL improves with small advantages when dealing with large corruption rates. SL outperforms ICL slightly across most corruption rates when using GPT2-XL. With multi-class classification, as the number of corrupted label increases, the performance of ICL gradually outperforms SL across all three models. The cause of different tendencies in performance between ICL and SL on binary and multi-class classification may be that binary classification tasks are generally simpler than multi-class tasks. With binary tasks, SL may be able to learn accurate decision boundaries more easily, while ICL may not provide additional benefit beyond the labeled examples. When faced with more complex multi-class tasks, ICL may better leverage the context of the inputs to better understand the relationships between the labels and provide more accurate predictions as the size of the model increases.

Discussion The learning behaviour of SL is strongly affected by the ground truth input-label correspondence as the model size increases. It suggests that SL requires correct pairings to perform downstream tasks. We hypothesize that, for ICL, the wrong pairings sabotage the pre-trained knowledge to some extent, but the large language models are more able to absorb knowledge directly from the pairings in the demonstrations. Similarly, input-label correspondence has a consistent effect on the learning behaviour of SL across all model sizes, but SL is more sensitive to corrupted labels compared to ICL. The reason for this may be that with SL, the model is fine-tuned by updating parameters using the wrong input-label pairs which enlarge the noisy “signal” in the labels. SL may be more prone to over-fitting than ICL, especially with high corruption rates. When the labels are noisy, the model may learn to fit the noise data during training rather than the underlying patterns, which may lead to

¹ Detailed performance for each dataset can be found in Technical Appendix A at https://github.com/xdwang0726/ICL_LL/blob/main/ECAI_Supplementary_Document.pdf.

² The detailed visualization can be found in Figure 1 in the Technical Appendix B at https://github.com/xdwang0726/ICL_LL/blob/main/ECAI_Supplementary_Document.pdf.

Dataset	LM	Noisy Level									
		0		25		50		75		100	
		SL	ICL	SL	ICL	SL	ICL	SL	ICL	SL	ICL
Binary	GPT2-Large	51.6 _{2.1}	53.9 _{16.9}	45.4 _{1.4}	50.4 _{11.2}	46.9 _{3.1}	44.5 _{3.7}	45.8 _{1.5}	40.9 _{2.9}	44.1 _{5.1}	41.1 _{1.8}
	GPT2-XL	49.4 _{2.5}	46.3 _{8.7}	45.0 _{2.2}	47.1 _{6.4}	46.6 _{1.0}	43.3 _{6.5}	45.7 _{1.6}	40.9 _{5.8}	43.0 _{4.9}	43.4 _{6.2}
	GPT-J	47.2 _{3.7}	54.0 _{13.4}	43.8 _{1.2}	48.0 _{13.7}	45.5 _{3.3}	43.5 _{3.6}	41.9 _{1.7}	38.6 _{2.3}	38.3 _{3.9}	37.8 _{1.7}
Multi-class	GPT2-Large	42.6 _{14.7}	21.9 _{12.1}	33.6 _{7.4}	22.3 _{13.6}	25.0 _{3.4}	19.5 _{12.4}	18.5 _{4.6}	18.8 _{13.1}	13.8 _{3.7}	17.7 _{11.0}
	GPT2-XL	41.9 _{14.0}	33.7 _{9.7}	33.2 _{8.1}	34.1 _{8.0}	24.5 _{2.9}	29.0 _{4.1}	20.3 _{3.2}	26.1 _{4.3}	13.0 _{2.8}	23.6 _{4.9}
	GPT-J	42.4 _{16.2}	40.6 _{11.4}	38.5 _{14.0}	36.0 _{10.2}	28.1 _{6.9}	29.2 _{3.5}	19.1 _{5.0}	23.3 _{6.8}	12.4 _{4.0}	20.9 _{7.9}
AVERAGE	GPT2-Large	47.1 _{10.6}	37.9 _{21.8}	39.5 _{8.0}	36.4 _{18.9}	35.9 _{12.4}	32.0 _{15.9}	32.2 _{15.2}	29.8 _{14.8}	28.9 _{17.1}	29.4 _{14.6}
	GPT2-XL	45.6 _{9.9}	40.0 _{10.7}	39.1 _{8.4}	40.6 _{9.6}	35.6 _{12.3}	36.2 _{9.2}	33.0 _{14.1}	33.5 _{9.3}	28.0 _{16.7}	33.5 _{11.9}
	GPT-J	44.8 _{10.9}	47.3 _{13.3}	41.2 _{9.4}	42.0 _{12.6}	36.8 _{10.7}	36.3 _{8.5}	30.5 _{12.9}	30.9 _{9.5}	25.4 _{14.6}	29.4 _{10.5}

Table 3. Performance comparison of supervised learning and in-context learning on noisy labels. We show the mean and standard deviation across different classification types. The subscripts are the standard deviations over 5 different seeds.

Dataset	LM	Imbalance Ratio					
		Low		Medium		High	
		SL	ICL	SL	ICL	SL	ICL
Binary	GPT2-Large	52.8 _{6.4}	49.1 _{11.3}	47.8 _{6.4}	49.4 _{10.8}	36.4 _{0.3}	46.1 _{11.9}
	GPT2-XL	51.7 _{6.5}	41.1 _{5.9}	46.4 _{4.4}	52.6 _{5.6}	35.3 _{1.5}	39.3 _{4.6}
	GPT-J	52.4 _{6.1}	52.5 _{12.2}	43.0 _{1.7}	53.6 _{17.0}	34.3 _{0.5}	53.5 _{20.2}
Multi-class	GPT2-Large	48.1 _{9.8}	25.6 _{15.4}	43.8 _{15.3}	23.1 _{15.8}	33.2 _{13.9}	21.8 _{11.1}
	GPT2-XL	44.6 _{5.6}	39.6 _{6.8}	47.0 _{12.8}	35.8 _{11.4}	38.5 _{14.8}	36.4 _{11.2}
	GPT-J	47.4 _{11.6}	48.7 _{17.5}	47.9 _{15.4}	51.8 _{17.3}	36.3 _{18.5}	38.2 _{12.2}
AVERAGE	GPT2-Large	50.5 _{7.8}	37.4 _{17.6}	45.8 _{10.7}	36.3 _{18.8}	34.8 _{8.9}	34.0 _{16.8}
	GPT2-XL	48.2 _{6.7}	40.4 _{5.8}	46.7 _{8.5}	44.2 _{12.2}	36.9 _{9.6}	37.9 _{7.8}
	GPT-J	49.9 _{8.7}	50.6 _{13.7}	45.4 _{10.2}	52.7 _{15.3}	35.3 _{11.8}	45.9 _{17.1}

Table 4. Performance comparison of supervised learning and in-context learning on label imbalance. We show the mean and standard deviation across different classification types. The subscripts are the standard deviations over 5 different seeds.

Model	β_1		β_0	
	SL	ICL	SL	ICL
GPT2-Large	-0.174	-0.094	0.454	0.378
GPT2-XL	-0.165	-0.080	0.445	0.407
GPT-J	-0.198	-0.187	0.456	0.465

Table 5. Sensitivity measure on the performance against the different noisy levels

poor generalization performances. Also, ICL takes into account the context in which the data are presented. In ICL, the model learns to predict the label based on not only the input features but also the context that can help to mitigate the impact of noisy and mislabeled data.

4.2 How do ICL and SL Learn with Imbalanced Labels?

Label Imbalance Matters Little to ICL As shown in Figure 4, label imbalance does not have significant effects on ICL. Specifically, decreasing the number of positive examples (i.e., increasing the level of label imbalance) does not lead to a significant drop in performance for any of the models. We also notice that learning with a low imbalance ratio (a balanced dataset) is not always the best choice, and a certain level of label imbalance may be expected or even desirable. For instance, by using GPT2-XL and GPT-J, the performances with medium imbalance ratios slightly outperform those with low imbalance ratios. One possible reason is that, in some cases, the presence of a reasonably high proportion of negative examples may actually help the language model learn to better discriminate between the positive and negative cases. This is because the language model can learn

to recognize common patterns or characteristics that are associated with the negative examples, which can then help it more easily identify the positive ones.

Model	β_1		β_0	
	SL	ICL	SL	ICL
GPT2-Large	-0.157	-0.034	0.515	0.376
GPT2-XL	-0.113	-0.025	0.495	0.420
GPT-J	-0.146	-0.047	0.508	0.520

Table 6. Sensitivity measure on the performance against the different imbalance ratio

Compared with SL The sensitivity measure in Table 6 shows that ICL is less sensitive to label imbalance compared to SL consistently across all three models. This indicates that SL is vulnerable to imbalanced labels, and the performance drops up to 15% from relatively balanced labels to highly imbalanced ones on average across three models. However, ICL is not so sensitive, with performance differences fluctuating between 3.4% and 6.8% over imbalance ratios. In addition, Figure 5 shows that the performance of ICL gradually surpasses SL as the model size increases.

In Table 4³, we also observe that, with binary classification, ICL shows sustained sensitivity as the imbalance ratio increases. ICL always outperforms supervised learning with medium and high imbalance ratios across all three models. With multi-class classification,

³ Detailed visualization can be found in Figure 2 in the Technical Appendix B at https://github.com/xdwang0726/ICL_LL/blob/main/ECAI_Supplementary_Document.pdf.



Figure 3. Performance comparison of supervised learning and in-context learning under different corrupted rates. We plot the mean accuracy (\pm standard deviation) across six different datasets for three different model sizes.




Figure 4. Results with varying imbalance ratios for in-context learning across three different models

though ICL is less sensitive to label imbalance, the model performance is not superior to that of SL, especially with GPT2-Large and GPT2-XL. With GPT-J, the advantage of ICL becomes obvious, which indicates that, with large language models, ICL is clearly preferable when dealing with class-imbalanced datasets.

Discussion In SL, imbalanced data imposes “label bias” during learning where the decision boundaries can be strongly influenced by the frequent classes. SL is incited to learn features from the frequent classes and may ignore other features. In contrast, ICL achieves non-trivial performance for classification tasks with imbalanced demonstrations as it may mitigate the effects of bias by leveraging the context specific information. This strongly suggests that the learning behaviour of ICL is not affected by the number of instances per class in the demonstrations, and ICL is capable of preserving pre-training knowledge as well as learning richer features. The reason for this may be that ICL can learn not only the label-relevant features but also other interesting features that capture the intrinsic properties of the input distributions and the structures of the input-label correspondence, which may generalize better to rare classes. This hypothesis needs further study.

5 Analysing ICL’s Sensitivity Using Attention Score

Our findings seem intuitive but are not trivial. The language models are not pre-trained on perturbed samples, but still ICL is less sensitive to label perturbations compared to SL. Understanding these observations are crucial for practitioners to apply ICL. The attention mechanism [2] has been used as an important component across LLMs. An attention layer produces a distribution over input representations to be attended to, and the distribution can reflect the significance of the inputs. Inspired by Wiegrefe and Pinter [38] that shows the utility of attention scores in interpretation, we further examine why ICL is less sensitive to label perturbations using attention scores. Specifically, we compute the attention scores of the output label word in each example for demonstration \mathcal{C} and sum up the scores for the same class together. We report the summed attention score ratio across different class labels (i.e. correct labels and perturbed labels) in the MR dataset and refer this as *Attention ratio score* that explains the robustness of ICL under the imbalance and noisy label settings.

LM	Ratios (corrected : corrupted)						STD
	5:0	4:1	3:2	2:3	1:4	0:5	
GPT2-Large	0.424	0.428	0.434	0.443	0.459	0.453	0.014
GPT2-XL	0.521	0.561	0.525	0.591	0.592	0.564	0.031
GPT-J	0.370	0.393	0.428	0.297	0.323	0.316	0.051

Table 7. Attention ratio scores for different correct label ratios for three models. Scores between models are not comparable, as attention scores vary. *STD* is the standard deviation across different corrupted ratios per model.

Noisy Labels To construct the experiment with noisy labels, we use 5 demonstration examples ($k = 5$) and evaluate with various corrupted rates. Results are shown in Table 7. With relatively small language models, such as GPT2-Large and GPT2-XL, there are small fluctuations of attention scores if more corrupted labels are added to the demonstration; however, with larger models like GPT-J, a large variation of attention appears as the number of corrupted labels increases. The aforementioned observation further indicates that smaller models cannot differentiate gold labels and corrupted ones and pay approximately the same attention to all labels; however, large language models demonstrate an ability to disambiguate between gold and corrupted labels, which supports our suspicion.




Figure 5. Performance comparison of supervised learning and in-context learning under different imbalance ratios. We plot the mean accuracy (\pm standard deviation) across six different datasets for three different model sizes.

Label Imbalance To construct the experiment with class-imbalanced data, we evaluate the attention score ratio of the two classes by fixing the number of examples of one class and adding examples for another. We report their attention score ratio in Table 8. The attention scores change slightly with different imbalance ratios, which indicates that each class receives approximately the same attention from ICL even with imbalanced label distributions.

LM	Imbalance Ratio						STD
	100%	50%	33%	25%	20%	10%	
GPT2-Large	0.946	0.956	0.936	0.967	0.913	0.887	0.030
GPT2-XL	0.929	0.995	0.983	0.995	0.913	0.975	0.035
GPT-J	0.909	0.825	0.847	0.857	0.908	0.913	0.038

Table 8. Attention ratio scores for different imbalance ratios for three models. Scores between models are not comparable, as attention scores vary. *STD* is the standard deviation across different imbalance ratios per model.

6 Related Work

6.1 Learning with Large Pre-trained Language Models

Large pre-trained language models have become a cornerstone of natural language processing, and have shown strong generalization abilities on a wide range of tasks [3, 9, 32, 33]. Currently, there are two key frameworks [23] to use LLMs, namely supervised fine-tuning and in-context learning.

Supervised learning performs gradient-based fine-tuning to adapt a task-agnostic pre-trained model for a specific task. Existing work [10, 17, 31] shows that supervised fine-tuning significantly boosts the performance on various NLP tasks, especially on small datasets.

In-context learning learns a new task without updating the model’s parameters. ICL adds a prompt (usually a task description) and annotated examples as demonstrations (also known as “shots”) to enable few-shot learning without fine-tuning [3]. Some recent work has been devoted to investigating ICL. For instance, Min *et al.* [27] tried to explain and understand the mechanism of ICL and found correct input-label mapping matters very little, while later work [41] further argued the impact of correct mappings under different configurations. Zhao *et al.* [45] showed that models are biased by the order of the demonstration examples. Liu *et al.* [24] and Mishra *et al.* [28] focused on the sensitivity of ICL: the former empirically studied the

sensitivity of GPT-3’s few-shot capabilities with respect to the selection of in-context examples and the latter re-framed the instructional prompts to improve the in-context performance for language models. Chen *et al.* [6] and Min *et al.* [26] used meta-training with an explicit in-context learning objective. Xie *et al.* [40] and Akyürek *et al.* [1] studied what enables ICL: the former explained ICL as implicit Bayesian inference and the latter used linear regression as a prototypical problem to study the learning algorithm of ICL. Yoo *et al.* [41] shared the same conclusion that ground-truth labels matter to ICL, as we claim, but by conducting experiments only on GPT-J. Our experiments cover three models with various sizes and we draw the more in-depth conclusion that LLMs are more sensitive to corrupted labels than smaller models. Parallel to our research, Zhang *et al.* [44] claimed that a well-balanced demonstration set does not consistently lead to better performance in ICL by studying the class imbalance under the demonstration example selection. While we reach a similar conclusion, we saw the result that label imbalance matters little to ICL when comparing ICL and SL.

6.2 Training with Perturbed Labels

Noisy Labels Learning with noisy labels can be divided into attempts that handle noisy labels in data pre-processing and attempts to learn robustly from the noisy labels directly. Some approaches preprocess auxiliary data by selecting clean instances according to rules. For instance, Li *et al.* [20] learned a teacher network with a small clean dataset to do importance re-weighting of the noisy labels in the loss function. Similarly, MentorNet [15] learned a data-driven curriculum to select clean instances that could guide the training of a student network. Further, Chen *et al.* [5] applied cross-validation to randomly split the noisy dataset and remove large-loss samples. Other approaches alleviate noise by making the learning process more robust to label noise, mainly among three categories: robust architectures [16, 18, 39] add a noise adaptation layer to learn label transitions or develop a dedicated architecture to support diverse types of label noise; robust regularization [25, 30] that enforces a deep neural network to overfit less to false-labeled examples explicitly or implicitly; and robust loss functions [12, 36] which adjust the loss according to the confidence of the labels.

Label Imbalance Imbalanced text classification focuses mainly on two types of methods: sampling-based and algorithm-level. Sampling-based methods balance class distributions by manipulating

training examples, including under-sampling, over-sampling, and a mix. Synthetic Minority Oversampling Technique (SMOTE) [4] is an over-sampling method where synthetic samples are generated for the minority class using the feature space to generate new instances with linear interpolation between positive instances. NearMiss [42] is an under-sampling technique that balances the class distribution by randomly eliminating majority class instances. However, over-sampling can potentially lead to over-fitting, while under-sampling may cause information loss. The algorithm-level approaches pursue balancing the class distribution and lifting the importance of minority class samples without altering the training data distribution, including cost-sensitive learning and threshold-adjustment methods. Cost-sensitive learning modifies cost functions [21, 37] and thresholding methods change the decision threshold at test time after training the classifier on the original imbalanced data [22].

7 Conclusion and Future Work

Our work is the first to study ICL among label perturbations. We empirically show that the learning behaviour of ICL is more influenced by the quality of the ground truth input-label pairings, and it is less effected by the number of pairings for each class in the demonstrations. In comparison with supervised learning (SL), we discover that ICL is less sensitive to perturbed labels than SL, which is still a standard method, and is still useful if the language model is relatively small. However, with large language models, the utility of ICL emerges. These findings highlight the importance of considering different learning strategies (ICL or SL) and model sizes when working with large language models and different classification tasks. Our work encourages practitioners to use ICL instead of SL or, at least, to consider evaluating the impact of perturbed pre-training on downstream tasks if large language models are available (here, we expected better performance from ICL and at least the model size should be larger than GPT-J). Our experiments are limited to text classification tasks, and we have not examined the performance of ICL for other tasks such as text generation nor in other domains such as computer vision. Extending our current work to such tasks is non-trivial, and is left for future work. To understand the learning behaviours of in-context learning and supervised learning, we train and fine-tune a model multiple times with different seeds on a fixed number of demonstration examples, which can be computationally intensive if the sizes of the datasets and models are large (for instance, we fine-tune the GPT-J model on the MR dataset using 4 NVIDIA A100 80G GPUs with parallel computing). This can be a non-trivial problem for practitioners with limited computational resources. Our experiments only have been conducted on the models that range from 774 million to 6.7 billion parameters due to these considerations. As yet, there does not seem to be any evidence that larger LLMs will reverse the tendencies we observed. We hope our findings of the behaviours of ICL under label perturbations can be a useful reference for potential users who have limited computational resources. We would also be interested to explore ways to optimize the learning behaviour of LLMs in different scenarios. We hope our work can inspire analysis of ICL in broader environments in the wild and provide insights for the design of future uses of ICL.

Acknowledgements

We would like to thank all reviewers for their comments, which helped improve this paper considerably. This research is partially

funded by The Natural Sciences and Engineering Research Council of Canada (NSERC) through a Discovery Grant to R. E. Mercer. F. Rudzicz is supported by a CIFAR Chair in AI.

References

- [1] Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou, ‘What learning algorithm is in-context learning? Investigations with linear models’, in *The Eleventh International Conference on Learning Representations (ICLR)*, (2023).
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, ‘Neural machine translation by jointly learning to align and translate’, in *3rd International Conference on Learning Representations (ICLR)*, (2015).
- [3] Tom Brown, Benjamin Mann, et al., ‘Language models are few-shot learners’, in *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901, (2020).
- [4] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer, ‘SMOTE: Synthetic minority over-sampling technique’, *Journal of Artificial Intelligence Research*, **16**(1), 321–357, (2002).
- [5] Pengfei Chen, Ben Ben Liao, Guangyong Chen, and Shengyu Zhang, ‘Understanding and utilizing deep neural networks trained with noisy labels’, in *Proceedings of the 36th International Conference on Machine Learning*, pp. 1062–1070, (2019).
- [6] Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis, and He He, ‘Meta-learning via language model in-context tuning’, in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 719–730, (2022).
- [7] Ido Dagan, Oren Glickman, and Bernardo Magnini, ‘The PASCAL Recognising Textual Entailment Challenge’, in *Proceedings of the First International Conference on Machine Learning Challenges: Evaluating Predictive Uncertainty Visual Object Classification, and Recognizing Textual Entailment*, pp. 177–190, (2005).
- [8] Marie-Catherine de Marneffe, Mandy Simons, and Judith Tonhauser, ‘The CommitmentBank: Investigating projection in naturally occurring discourse’, *Proceedings of Sinn und Bedeutung*, **23**(2), 107–124, (2019).
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, ‘BERT: Pre-training of deep bidirectional transformers for language understanding’, in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, (2019).
- [10] Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith, ‘Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping’, *arXiv preprint: arXiv:2002.06305*, (2020).
- [11] Shivam Garg, Dimitris Tsipras, Percy Liang, and Gregory Valiant, ‘What can transformers learn in-context? a case study of simple function classes’, in *Advances in Neural Information Processing Systems*, volume 35, pp. 30583–30598, (2022).
- [12] Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel, ‘Using trusted data to train deep networks on labels corrupted by severe noise’, in *Advances in Neural Information Processing Systems*, volume 31, pp. 10456–10465, (2018).
- [13] Eduard Hovy, Laurie Gerber, Ulf Hermjakob, Chin-Yew Lin, and Deepak Ravichandran, ‘Toward semantics-based answer pinpointing’, in *Proceedings of the First International Conference on Human Language Technology Research*, pp. 1–7, (2001).
- [14] Maksims Ivanovs, Roberts Kadikis, and Kaspars Ozols, ‘Perturbation-based methods for explaining deep neural networks: A survey’, *Pattern Recognition Letters*, **150**, 228–234, (2021).
- [15] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei, ‘MentorNet: Learning data-driven curriculum for very deep neural networks on corrupted labels’, in *Proceedings of the 35th International Conference on Machine Learning*, pp. 2304–2313, (2018).
- [16] Ishan Jindal, Matthew Nokleby, and Xuewen Chen, ‘Learning deep networks from noisy labels with dropout regularization’, in *2016 IEEE 16th Int. Conference on Data Mining (ICDM)*, pp. 967–972, (2016).
- [17] Cheolhyoung Lee, Kyunghyun Cho, and Wanmo Kang, ‘Mixout: Effective regularization to finetune large-scale pretrained language models’, in *Proceedings of International Conference on Learning Representations (ICLR)*, (2020).
- [18] Kimin Lee, Sukmin Yun, Kibok Lee, Honglak Lee, Bo Li, and Jinwoo

- Shin, 'Robust inference via generative classifiers for handling noisy labels', in *Proceedings of the 36th International Conference on Machine Learning*, pp. 3763–3772, (2019).
- [19] Xin Li and Dan Roth, 'Learning question classifiers', in *COLING 2002: The 19th Int. Conference on Computational Linguistics*, (2002).
- [20] Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li, 'Learning from noisy labels with distillation', in *2017 IEEE Int. Conf. on Computer Vision (ICCV)*, pp. 1928–1936, (2017).
- [21] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár, 'Focal loss for dense object detection', in *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2999–3007, (2017).
- [22] Charles X. Ling and Victor S. Sheng, 'Cost-sensitive learning', in *Encyclopedia of Machine Learning*, eds., Claude Sammut and Geoffrey I. Webb, 231–235, Springer, (2008).
- [23] Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin Raffel, 'Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning', in *Advances in Neural Information Processing Systems*, volume 35, pp. 1950–1965, (2022).
- [24] Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen, 'What makes good in-context examples for GPT-3?', in *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pp. 100–114, (2022).
- [25] Aditya Krishna Menon, Ankit Singh Rawat, Sashank J. Reddi, and Sanjiv Kumar, 'Can gradient clipping mitigate label noise?', in *Eighth International Conference on Learning Representations (ICLR)*, (2020).
- [26] Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi, 'MetaICL: Learning to learn in context', in *Proc. of the 2022 Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2791–2809, (2022).
- [27] Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer, 'Rethinking the role of demonstrations: What makes in-context learning work?', in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 11048–11064, (2022).
- [28] Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, and Hannaneh Hajishirzi, 'Reframing instructional prompts to GPTk's language', in *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 589–612, (2022).
- [29] Bo Pang and Lillian Lee, 'Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales', in *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, pp. 115–124, (2005).
- [30] Gabriel Pereyra, G. Tucker, Jan Chorowski, Lukasz Kaiser, and Geoffrey E. Hinton, 'Regularizing neural networks by penalizing confident output distributions', in *5th International Conference on Learning Representations, (ICLR)*, (2017).
- [31] Jason Phang, Thibault Févry, and Samuel R. Bowman, 'Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks', *arXiv preprint arXiv:1811.01088*, (2018).
- [32] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al., 'Language models are unsupervised multitask learners', *OpenAI blog*, **1**(8), 9, (2019).
- [33] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu, 'Exploring the limits of transfer learning with a unified text-to-text transformer', *J. of Machine Learning Research*, **21**(140), 1–67, (2020).
- [34] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts, 'Recursive deep models for semantic compositionality over a sentiment treebank', in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, (2013).
- [35] Ben Wang and Aran Komatsuzaki, GPT-J-6B: A 6 billion parameter autoregressive language model. <https://github.com/kingoflolz/mesh-transformer-jax>, May 2021.
- [36] Ruxin Wang, Tongliang Liu, and Dacheng Tao, 'Multiclass learning with partially corrupted labels', *IEEE Transactions on Neural Networks and Learning Systems*, **29**(6), 2568–2580, (2018).
- [37] Shoujin Wang, Wei Liu, Jia Wu, Longbing Cao, Qinxue Meng, and Paul J. Kennedy, 'Training deep neural networks on imbalanced data sets', in *2016 International Joint Conference on Neural Networks (IJCNN)*, pp. 4368–4374, (2016).
- [38] Sarah Wiegrefe and Yuval Pinter, 'Attention is not not explanation', 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)*, pp. 11–20, (2019).
- [39] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang, 'Learning from massive noisy labeled data for image classification', in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2691–2699, (2015).
- [40] Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma, 'An explanation of in-context learning as implicit bayesian inference', in *The Tenth Int. Conf. on Learning Representations, (ICLR)*, (2022).
- [41] Kang Min Yoo, Junyeob Kim, Hyuhng Joon Kim, Hyunsoo Cho, Hwiyeol Jo, Sang-Woo Lee, Sang-goo Lee, and Taeuk Kim, 'Ground-truth labels matter: A deeper look into input-label demonstrations', in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 2422–2437, (2022).
- [42] J. Zhang and I. Mani, 'KNN approach to unbalanced data distributions: A case study involving information extraction', in *Proceedings of the ICML'2003 Workshop on Learning from Imbalanced Datasets*, (2003).
- [43] Xiang Zhang, Junbo Zhao, and Yann LeCun, 'Character-level convolutional networks for text classification', in *Advances in Neural Information Processing Systems*, volume 28, pp. 649–657, (2015).
- [44] Yiming Zhang, Shi Feng, and Chenhao Tan, 'Active example selection for in-context learning', in *Proc. of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 9134–9148, (2022).
- [45] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh, 'Calibrate before use: Improving few-shot performance of language models', in *Proceedings of the 38th International Conference on Machine Learning*, pp. 12697–12706, (2021).